

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
21 February 2002 (21.02.2002)

PCT

(10) International Publication Number  
**WO 02/15494 A1**

- (51) International Patent Classification<sup>7</sup>: **H04L 12/56, H04J 3/22, 3/24**
- (21) International Application Number: **PCT/US01/24855**
- (22) International Filing Date: **7 August 2001 (07.08.2001)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:  
09/635,695 10 August 2000 (10.08.2000) US  
09/652,140 31 August 2000 (31.08.2000) US
- (71) Applicant: **VERIZON COMMUNICATIONS INC.**  
[US/US]; 1095 Avenue of the Americas, New York, NY 10036 (US).
- (72) Inventors: **VOIT, Eric, A.; 5611 Oakmont Avenue, Bethesda, MD 20817 (US). BAUM, Robert, T.; 429 Girard Street #304, Gaithersburg, MD 20877 (US).**
- (74) Agent: **SUCHYTA, Leonard, C.; c/o Christian R. Andersen, 600 Hidden Ridge Drive, Mailcode HQE03H01, Irving, TX 75038 (US).**
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European Patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

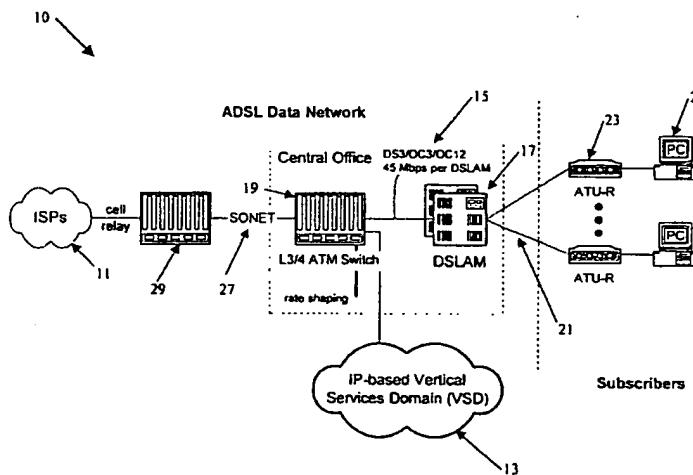
— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: AUTOMATED SERVICE PROVISIONING IN COMBINATION OF VERTICAL SERVICES AND DIGITAL SUBSCRIBER LINE DOMAINS



**WO 02/15494 A1**



(57) Abstract: The provided network architecture introduces QoS into the ADN, in a manner that enables the delivery of sophisticated and demanding IP-based services (13) to subscribers (25), does not affect existing Internet tiers of service, and is cost-effective in term of initial costs, build-out, and ongoing operations. The architecture utilizes a switch (19) capable of examining and selectively forwarding packets or frames based on higher layer information in the protocol stack, that is to say on information that is encapsulated in the layer-2 information utilized to define normal connectivity through the network. The switch (19) enables segregation of upstream traffic by type and downstream aggregation of Internet traffic together with traffic from a local vertical services domain (13). Systems coupled to the network and software in a user's computer enable a method for automated, end-to-end provisioning of a logical data circuit, in response to a customer's request for bundled broadband services.

AUTOMATED SERVICE PROVISIONING IN  
COMBINATION OF VERTICAL SERVICES  
AND DIGITAL SUBSCRIBER LINE DOMAINS

Technical Field

5        Certain concepts involved in the present invention relate to techniques for implementing data communication services, for example in a local access network utilizing digital subscriber line technology, to support quality of service (QoS) and local introduction of vertical services. Other concepts involved in the present invention relate to techniques for automatically provisioning services  
10      through such a network.

Background

Modern society continues to create exponentially increasing demands for digital information and the communication of such information between data devices. Local area networks use a network, cable or other media to link stations  
15      on the network for exchange of information in the form of packets of digital data. These networks have proven quite successful in providing data communications in commercial applications. However, the common local area network architectures require installation of specialized wiring and use of specific wiring topologies. For example, the most popular network protocols, such as Ethernet,  
20      require special rules for the wiring, for example with regard to quality of wire, range of transmission and termination. Furthermore, to extend communications to a wider domain still requires connection of at least one node of the local area network out to a wider area network, such as the network of an Internet Service Provider (ISP). High speed links enabling such wide area access from a LAN  
25      domain, for example using T1 lines, are quite expensive and justified only for hi-end commercial users.

The most common form of computer-to-computer communication in use today, particularly for wide area communications, still relies on modems and analog telephone network connections. The telephone-based operation provides the voice grade analog modem a unique power, the necessary connections are virtually ubiquitous. Such modems can communicate via almost any telephone line or wireless telephone (e.g. cellular) to any other such telephone connection, virtually anywhere in the world. The telephone network, however, was designed to provide approximately 3.3 kHz of analog voice bandwidth. Consequently, the data rates that are possible through the telephone network are quite low. Even with a variety of recent enhancements, the data speeds remain at or below 56 kbps.

Integrated Services Digital Network (ISDN) offers somewhat faster data communications and the capacity for concurrent data and voice telephone services. The 160 kb/s capacity carries two bearer (B) channels, each at 64 kb/s, one data (D) channel at 16 kb/s and overhead information in a 16 kb/s embedded operations channel (EOC). The two B-channels may be used separately, for example, for one voice telephone call and one data communication session. The D-channel typically is used for signaling, for call set-up and the like. Some applications allow aggregation of the channels, to combine the B-channels and possibly the D-channel to provide data communications up to the combined rate of 144 kb/s. However, these data rates offered by ISDN already are too slow for many multimedia applications. The high-speed and wide availability of modern personal computers (PCs) continually gives rise to ever more sophisticated multimedia applications. Communications for such applications, typically between the PC and the Internet, already are driving the need for speed to rates far above those available on normal ISDN lines.

A number of technologies are being developed and are in early stages of deployment, for providing substantially higher rates of data communication, for example ranging from 640 kb/s to 7.1 Mb/s. For example, cable television companies are now beginning to offer 'cable modem' services, which allow customers to communicate data over available bandwidth on the coaxial cable of a cable television network. After considering several other options, a number of

- 3 -

the local telephone carriers are working on enhancements to their existing copper-wire loop networks, based on various xDSL technologies.

The term xDSL here is used as a generic term for a group of higher-rate digital subscriber line communication schemes capable of utilizing twisted pair wiring from an office or other terminal node of a telephone network to the subscriber premises. Examples under various stages of development include ADSL (Asymmetrical Digital Subscriber Line), HDSL (High data rate Digital Subscriber Line) and VDSL (Very high data rate Digital Subscriber Line).

The telephone carriers originally proposed use of ADSL and similar high-speed technologies to implement digital video services, for example in networks sometimes referred to as video 'dialtone' networks. The ADSL line technology provided a mechanism for high-speed transport of MPEG encoded video information to video terminal devices in the customers' homes. Examples of such ADSL-based video dialtone networks are disclosed in U.S. Patent Nos. 5,247,347, 5,410,343 and 5,621,728. The carriers are now deploying a range of xDSL data services targeted at high-speed Internet access and high-speed access to private data networks. U.S. Patent No. 5,790,548 to Sistanizadeh et al. discloses an example of an ADSL based data network, e.g. for high-speed access to the Internet and to corporate LANs.

The current design goals of DSL data networks for Internet access do not support high-end vertical services, that is to say services demanding IP-based applications that require assurance of some level of quality of service (QoS). For example, packet-switched Voice over IP (VoIP) requires low latency, low jitter (i.e., a relatively constant bit rate), and non-correlated packet loss. Streaming video has similar requirements, and in addition, requires high bandwidth. DSL data networks designed to support high speed Internet and Intranet access have been optimized to support traffic that is bursty and is not sensitive to latency or jitter. For example, current implementations supporting ATM cell traffic employ the Unspecified Bit Rate (UBR) class of service, which does not provide any bandwidth or delay guarantees. Consequently, transport of video materials through such DSL data networks inflicts video delays, loss of audio/video synchronization, and image fragmentation.

- 4 -

Furthermore, lengthy bandwidth intensive sessions for video or other broadband applications may degrade the throughput to all other subscribers served through a shared node, such as a gateway router or a concentrated link. For two-way video, upstream will have even worse quality and throughput problems, due to the best effort nature of the DSL data network implemented for Internet access and because the upstream bandwidth is significantly less than that of the downstream channel.

To appreciate the situation and problems, it may be helpful here to consider an ADSL data implementation of a local access network, as a representative example, in somewhat more detail. Fig. 9 is a block diagram of a typical ADSL data network of the type currently in-use by a number of incumbent and competitive local exchange carriers to provide high-speed access to Internet Service Providers (ISPs) and thus to the Internet. Fig. 10 provides an alternative functional illustration of the elements of such a network. Of particular note, Fig. 15 shows the various protocol stacks in association with the appropriate network elements.

As shown in Fig. 9, a central office (CO) 100 provides plain old telephone service (POTS) and digital subscriber line data service for a number of customers. For purposes of discussion, assume that the equipment at each of the various 20 customer premises 200 connects directly to the CO 100 via twisted pair type copper wiring 300. In an actual implementation, many customers may connect through such wiring to a remote terminal linked to the CO via optical fiber.

At each customer premises 200 in our example, the copper loop 300 carrying both the POTS and ADSL signals connects through a Network Interface Device (NID) 201 placed at the side of the home. A two pair loop is installed from the NID to the location where the ADSL unit 203, typically an ATU-R modem, is located in the home. One pair connects all of the signals on the line 300 from the NID 201 to the ADSL modem 203. Within the ATU-R type modem 203, a passive splitter/combiner type filter segregates the POTS signal and the data signals. The 25 POTS signal is transmitted over the second twisted pair back to the NID 201. The POTS line is then connected to the in-home wiring extensions at the NID 201, for distribution to one or more standard telephone devices 205 in the home.

- 5 -

Within the ATU-R type ADSL modem 203, the downstream coded ADSL signal is demodulated and decoded to an appropriate data interface protocol for connection to the PC 215. The PC 215 or other data device (Fig. 10) also sends data to the ADSL modem 203. The modem 203 modulates the upstream data and transmits appropriate signals over the line 300<sub>1</sub> or 300<sub>2</sub> to the corresponding modem 113<sub>1</sub> or 113<sub>2</sub> in the CO 100 (Fig. 9). The ATU-R interface may support bridging, such that multiple users can share the ADSL modem 203, for two-way data communication through the CO 100.

The lines 300 for the customer premises 200 connect through the main distribution frame (MDF) 101 to a Digital Subscriber Line Access Multiplexer (DSLAM) 111. The DSLAM includes a bank of ADSL terminal units of the type intended for central office applications, identified as ATU-Cs 113. The DSLAM also includes a multiplexer/demultiplexer (MUX) 115.

Within the DSLAM 111, each customer line 300 connects to an assigned ADSL terminal unit 113 in the central office (ATU-C). In the example illustrated, the first customer's line 300<sub>1</sub> connects through the MDF 101 to a first ATU-C 113<sub>1</sub> in the CO 100. The second customer's line 300<sub>2</sub> connects through the MDF 101 to a second ATU-C 113<sub>2</sub> in the CO 100. The ATU-C type ADSL units 113 include appropriate frequency dependent combiner/splitters, for segregating out the voice telephone traffic. Thus each ADSL unit 113 provides a connection for telephone traffic from the associated line 300 to the POTS switch 103.

The ADSL units 113 in the CO (ATU-Cs) essentially act as modulator/demodulators (modems) for sending and receiving data over the subscriber telephone lines 300. On the network side, each of the ATU-Cs 113 connects to the MUX 115. The MUX 115 multiplexes and demultiplexes the upstream and downstream data for the ADSL modems 113 and provides a connection to a high-speed link 119. Through subtending, the MUX 115 may also provide a data concentration for the communications over the link 119.

In a typical implementation, the concentrated data communications utilize a DS-3 link 119. However, because of increasing traffic demands, it is becoming necessary to upgrade the link 119 to SONET optical fiber, such as OC-3 or in some cases even OC-12. The link 119 provides two-way data communication between

the central office 100 and a data hub 121. In practice, this is a relatively long or wide area link using expensive interoffice facilities.

On the upstream side, the high-speed interoffice link 119 terminates on an ATM switch 123 for the ADSL data network (ADN). Although only one link 119 appears in the drawing, the asynchronous transfer mode (ATM) switch 123 will typically service a number of DSLAMs 111 in various end offices via similar DS or OC links. The ATM switch 123, in turn, provides a high-speed connection to a gateway router 125 coupled to an ATM cell relay network 129. Typically, the ATM switch 123 will aggregate traffic from a number of such links 119 onto an OC-3 or higher rate SONET link to the router 125. The router 125 and the cell relay network 129 enable transport of ATM cells for the subscribers to and from equipment of one or more Internet Service Providers (ISPs), shown by way of example as a concentrator 131 coupled to the public packet switched network commonly known as the Internet 132.

The illustrated local access type ADN network provides ATM cell transport from a customer premises 200 to the ISP concentrator 131. The ATM cells serve as the layer-2 routing or switching protocol for the lowest level definition of connectivity between two points of the network. Higher level protocols ride within the ATM cells.

The ATU-Rs 203 and the customer premises data equipment 215 connect via an Ethernet coupler. The customers' equipment communicates across the ADSL data network utilizing Ethernet, and the wide area communication involves transport of Internet protocol information typically in TCP/IP frames within Ethernet frames. The Ethernet frames carrying the TCP/IP frames are adapted into ATM cells. Attention is directed to the protocol stacks illustrated in the lower half of Fig. 10.

To efficiently provide cell relay, each customer is assigned an ATM virtual circuit that extends from the ATU-R 203 in the respective customer premises 200 to the gateway router 125. Although it was originally envisioned that ATM would support switched logical channels or virtual circuits, to date, such logical switching has proven impractical to implement and administer. Consequently, current practical ATM networks actually utilize permanent virtual circuits, not

- switched virtual circuits. For a given subscriber, the carrier therefore provisions an ATM permanent virtual circuit from the ATU-R 203 to the gateway router 125. The carrier programs one or more nodes along the path of that logical circuit, particularly the DSLAM 111, to regulate traffic on the virtual circuit to the upstream and downstream rates corresponding to the grade of service to which the particular customer subscribers. All data traffic for the subscriber goes over the entire length of the permanent virtual circuit, and most if not all nodes along that path limit that traffic to the rates of the subscription as defined in the provisioning data.
- The virtual circuit may be thought of as a solid pipe. All traffic passes through the entire length of the pipe-like virtual circuit, regardless of how many switches or other nodes the circuit passes through. The layer-2 protocol defining the circuit carries all of the higher level traffic end-to-end. Higher layer protocols are visible only at the ends of the pipe. Hence, any traffic flow processing intended to utilize the higher layers must occur at some point past one end or the other end of the virtual circuit.

The gateway router 125 also terminates permanent virtual circuits through the cell relay network 129 going to/from the ISP concentrators 131. The gateway router 125 aggregates traffic between a number of subscribers and each respective ISP. The ISP equipment 131 typically implements a variation of a point-to-point protocol (PPP) specifically adapted to ride over Ethernet, referred to as "PPP over Ethernet" (PPPoE). The virtual circuits to the ISPs, however, do not have sufficient capacity to simultaneously carry all subscriber traffic at the maximum rates of the customers' subscriptions. The MUX 115, the ATM switch 123, and the gateway router 125 concentrate and regulate the subscriber traffic going to and from the ISPs, typically on some type of "best efforts" basis.

In a typical Internet access service offering, the most expensive service tier provides 7.1 Mbps for downstream communication and 680 kbps for upstream communication. The next grade of service provides 1.6 Mbps for downstream communication and 90 kbps for upstream communication, whereas the lowest tier of service provides 640 kbps for downstream communication and 90 kbps for upstream communication. The maximum grade of service offered to an

individual subscriber depends on the rates for which the subscriber's line can qualify, although the subscriber may opt for a lower rate service since the higher-rate service is more expensive.

The approach outlined above relative to Figs. 9 and 10 works well for 5 Internet access if the traffic relates to web access, file transfers and the like, which do not require guaranteed quality of service. Various segments of the Internet industry, however, are rapidly developing new multimedia services and applications that already are pushing the capabilities of such a network. For example, increasingly, Internet traffic includes a number of types of 10 communication that require a guaranteed quality of service. Voice telephone communication over IP is extremely sensitive to latency and jitter. The permanent virtual circuits provide an unspecified bit rate (UBR) service and do not guarantee any minimal amount of delay or jitter. Also, because the rates are set by subscription, the service tends to be relatively inflexible. Some services, 15 such as multicasting of broadband information from the Internet into the local access ADN for a large number of concurrent users, can quickly overload one or more nodes or critical links of the network, for example the link 119 between the DSLAM 111 and the ATM switch 123 at the hub 121.

Most industry experts propose to increase the services available via the 20 public Internet. However, because the higher layer protocols are visible only on the Internet side of the virtual circuit "pipe," these services all must be implemented out past the end of the virtual circuit, at least behind the gateway router 129 and most likely in the public network, where it is possible to view and route based on higher level protocols, particularly Internet protocol (IP). Such a 25 migration strategy to implement new services creates severe problems. For example, in the network of Fig. 9, if a customer at premises 200<sub>1</sub> desired to order a video on demand, the customer would communicate via the assigned permanent virtual circuit and the ISP to a server on the Internet 132. The server would send the video stream back through the Internet 132, the ISP equipment 30 131, the cell relay network 129 and the virtual circuit from the router 125 to the ATU-R 203 for handoff to a PC or the like at 215. If the rate of the requested video exceeds the customer's subscription rate, the customer could not view the

video in real time during the download. Even if the rate of the requested video is below the customer's subscription rate, loading in the Internet or the local access network may impose delays and/or jitter in communication of some segments of the requested video. Assuming that the hub 121 and the links 119 implement a 5 subscriber concentration, ordering of videos or similar broadband files from the Internet 132 quickly consumes the shared resources through the hub 121 and the links 119, reducing the rates of service provided to other customers seeking concurrent Internet access.

It might be possible to increase the capacity of the links 119 and/or the 10 hubs 121; however, this tends to increase the carrier's recurring costs and often makes the overall service(s) of the ADN network economically impractical.

It has also been suggested to provide customers guaranteed quality of services for some portion of their communications, by segregating the traffic carried between the customer premises and the hub 121. This would require 15 assigning a plurality of ATM permanent virtual circuits to each subscriber, one for each different guaranteed level of quality of service and one for all other Internet traffic for the subscriber. Administration and provisioning of one virtual circuit per subscriber is already complicated, and the number of virtual circuits through any given ATM node is limited by current equipment designs. 20 Expanding the number of permanent virtual circuits per subscriber to support multiple QoS tiers of service therefore would be quite expensive, and the management thereof would become a nightmare. To support an increased number of virtual circuits, many having guaranteed QoS requiring some substantial minimum rate at all times, would also require that the operator 25 substantially upgrade the network to increase the end-to-end capacity all the way to the wide area network 132.

Furthermore, to actually receive the desired QoS requires that all elements involved in the communication must guarantee the desired level or quality of service. For communications across the public Internet 132, this means that 30 various nodes and links on the public Internet must be available and capable of providing a guarantee of the desired QoS. In point of fact, few nodes on the public Internet actually support any type of QoS. Hence, even if the ADN

- 10 -

supported a desired QoS, most subscribers would not benefit from that service because their communications over the public Internet would have no QoS guarantee, and would suffer from the usual problems of latency and jitter.

Consequently, current deployments of ADSL-based data networks, such as  
5 shown in Figs. 9 and 10 generate many customer complaints. From the customer perspective, the service does not deliver the data rates that the customer pays for on a consistent basis. The customer typically blames such problems on network equipment failure. In fact, most of the problems already are due to virtual circuit congestion problems, of the kinds outlined above. Essentially, the ADN network  
10 is crippled by the unpredictable nature of the service levels that the customers perceive due to congestion on the ADN and on the public Internet.

Also, with this approach, because all of the major service elements are implemented in servers accessible to the Internet, all of the services are subject to severe security risks. Each service provider's server is accessible to virtually any  
15 computer coupled for communication via the Internet. This openness is a desirable feature of the public Internet. However, a consequence is that any such server is accessible to and thus subject to attack from any hacker having Internet communications capabilities. Popular services, particularly those generating substantial revenues, become prime targets for attack.

20 Another area of problems is that the ADN does not offer the carrier any technique for offering its own differentiated service applications. To compete with other service providers, the carrier operating the ADSL-based data network needs to introduce its own multimedia services, for example, its own video services to compete with video services of cable television companies (that offer  
25 competing Internet access services). As noted above, however, introduction of a new service, such as true video on demand or broadcast video requires communications via the public Internet 132. This is true even if the carrier operating the network of Figs. 9 and 10 wanted to initiate its own video service(s).

30 Hence, there is an ongoing need to improve the architecture and operation of a digital subscriber line data communication network, particularly to facilitate finer gradation of services within the local network. The need, first, is for such a

local network to support introduction of services on a 'vertical' basis within the local access network separate and apart from the common forms of Internet traffic, both for commercial differentiation and for increased security. In a related need, the local network needs to support a number of different levels of  
5 quality of service (QoS).

### Summary

A general objective of the invention is to implement an enhanced digital communication network for subscriber lines that supports vertical introduction of new communication and/or multimedia services.

10 A further objective is to support multiple levels or grades of quality of service within the access network.

Another objective of the invention relates to improvement of the cost effectiveness of the data network, for example, by reducing the demand for high-capacity interoffice links while increasing the bandwidth available at the network  
15 edge for at least some types of services.

A related objective is to provide a technique for introduction of new high-end services near the network edge, from a domain that is more secure and therefore less subject to hacker attacks.

20 A further objective of the invention is to support QoS and/or local introduction of vertical services, without the need to assign multiple virtual circuits or the like to each subscriber.

Still further objectives of the invention relate to provisioning of service through an access data network. Specifically, it is an objective of the provisioning related concepts to significantly reduce, or completely remove, the errors present  
25 in the existing data circuit provisioning process, which includes: manual ordering processes; manual record-keeping; and manual circuit changes at the PSTN frame and the intermediate ADSL frame.

Aspects of the invention relate to unique methods and network architectures for providing a combination of wide area internetwork service and

vertical communication services via a local access network. Other aspects of the invention relate to a particular switch developed to facilitate a unique form of routing, for example to support of QoS and vertical service insertion.

Hence, a first aspect of the invention relates to a method of segregating traffic for at least two different network domains. The method involves establishing a contiguous layer 2 protocol connectivity, upstream from a customer premises to a communication access node coupled to a first network domain. At an intermediate point along the contiguous connectivity, at least two types of upstream transmissions from the customer premises are distinguished, based on information encapsulated within the layer-2 protocol. The methodology forwards each distinguished transmission of a first transmission type from the intermediate point to the first network domain. Similarly, the method entails forwarding each detected transmission of a different second type from the intermediate point to a second network domain, which is logically separated from the first network domain.

Another method aspect of the invention involves providing a combination of access to two different network domains through an access data network. In this case, the actual method includes a step of provisioning a logical communication circuit extending from a customer premises through the access data network to a communication access node coupled to a first network domain. This provisioning comprises defining the logical communication circuit in terms of a layer-2 protocol. An intermediate node along the logical communication circuit examines communicated information in transmissions from the customer premises, for a protocol encapsulated within the layer-2 protocol, in order to distinguish transmission types. The method also includes steps of forwarding transmissions of two different types. A first type of detected transmission is forwarded from the intermediate node to the communication access node, over the logical communication circuit defined in terms of the layer-2 protocol. A second type of detected transmission is forwarded to a second network domain logically separate from the first network domain.

Another aspect of the invention relates to a method involving providing rate adaptive digital subscriber line communications from a digital subscriber

- 13 -

line access multiplexer to a plurality of subscribers, over respective subscriber lines. Consequently, the communications can operate at maximum rates that respective line conditions will allow. For each subscriber, wide area access services are provided from an access router node coupled to the digital subscriber line access multiplexer to a node coupled to a wide area internetwork. The rates for the access services conform to service level agreements with respective subscribers. For each subscriber, the method also entails aggregating and segregating communications for a vertical services domain with those for the wide area access services, for combined communication via the maximized-rate digital subscriber line communications.

From a somewhat different perspective, another inventive method offers a combination of wide area internetwork service and vertical communication services through a local access network. This method comprises provisioning a logical communication circuit to support a subscriber-selected grade of service to the wide area internetwork. The logical communication circuit extends from a customer premises of the subscriber to a communication access node coupled to the wide area internetwork. The provisioning involves defining the logical communication circuit in terms of a layer-2 protocol, for example in terms of an ATM virtual circuit. An intermediate node along the communication circuit examines communicated information in transmissions from the customer premises, for protocol layers higher than the layer-2 protocol. In this manner, the intermediate node distinguishes transmission types. The inventive method entails routing each detected transmission of a first type from the intermediate node, to the communication access node, over the logical communication circuit defined in terms of the layer-2 protocol, at a transmission rate corresponding to the subscriber selected grade of service. Each detected transmission of another type is extracted from the logical communication circuit at the intermediate node. The methodology routes each extracted transmission to a vertical services network coupled locally to the intermediate node.

Another feature of this later methodology relates to the processing of downstream communications. Here, the method further comprises receiving a first downstream transmission intended for the subscriber from the

communication access node, at the intermediate node. This transmission is received over the logical communication circuit, at a downstream transmission rate corresponding to the subscriber-selected grade of service. The intermediate node also receives a second downstream transmission intended for the  
5 subscriber, from the vertical services network. The method entails inserting the second downstream transmission into the logical communication circuit, to combine the first and second downstream transmissions for communication over the logical communication circuit from the intermediate node to the customer premises.

10 Other inventive aspects embody a data communication network, for example, for providing a combination of wide area internetwork service and vertical communication services. The network comprises a communication access node coupled to a first network domain, such as the wide area internetwork. A physical communication circuit provides a coupling to customer  
15 premises communication equipment. The network also includes a logical communication circuit, for example provisioned to support a subscriber-selected grade of service to the first network domain. The logical communication circuit extends through the physical communication circuit and through the network, to the communication access node coupled to the wide area internetwork. The  
20 provisioning defines the logical communication circuit in terms of a layer-2 protocol. The network further comprises an intermediate communication node, coupled to the physical communication circuit at an intermediate location along the logical communication circuit. The intermediate communication node includes a controller for examining communicated information in transmissions  
25 from the customer premises, for one or more protocol layers encapsulated with the layer-2 protocol. This higher level examination enables the intermediate node to distinguish transmissions by type, for example by different Ethertype protocols carried in layer-2 type ATM cells. The intermediate communication node also includes a switch fabric, such as a router or the like, coupled to the  
30 logical communication circuit and responsive to the controller. This logical switch fabric enables forwarding of each detected transmission of a first transmission type from the intermediate node to the communication access node over the logical communication circuit. The transmissions forwarded to the

communication access node preferably travel at a transmission rate corresponding to the subscriber-selected grade of service. The logical switch fabric also enables forwarding of each detected transmission of any other type from the logical communication circuit to a second network domain, such as the  
5 vertical services network.

From a somewhat different perspective, another inventive network aspect relates to a digital subscriber line data network. This data network provides a combination of wide area internetwork service and vertical communication services. The data network comprises a communication access node coupled to  
10 the wide area internetwork. Digital subscriber line transceivers, coupled to network ends of subscriber lines, provide data communication with transceivers coupled to customer premises ends of respective subscriber lines. Disclosed examples of the transceivers of the network are ATU-C units. The data network further comprises an access switch coupled to the digital subscriber line  
15 transceivers, for example through a mutliplexer. The access switch receives data from customer premises equipment via respective digital subscriber line transceivers, and it supplies data intended for transmission to predetermined customer premises equipment to the respective digital subscriber line transceivers. The data network includes a high-speed data link between the  
20 access switch and the communication access node.

The network also includes a logical communication circuit provisioned in terms of a layer-2 protocol for each subscriber line. The logical circuit extends through the switch and the high-speed data link. In the preferred embodiments, each logical communication circuit is provisioned to support a respective  
25 subscriber-selected grade of service to the wide area internetwork and to extend from a respective customer premises to the communication access node. The inventive data network also comprises a vertical services network coupled locally to the access switch. A controller examines communicated information in upstream transmissions on each logical communication circuit, for protocol  
30 layers higher than the layer-2 protocol, to distinguish transmission types.

In operation, the access switch routes each detected transmission of a first transmission type, received from a customer premises over the respective logical

communication circuit on the respective line, over the respective logical communication circuit on the high-speed data link to the communication access node. In the preferred embodiments, such communications travel at a transmission rate corresponding to the respective subscriber-selected grade of service. The access switch also forwards each detected transmission of any other type from the respective logical communication circuit to the vertical services network.

A still further inventive aspect relates to an implementation of a data switch for use in a communication network. The data switch comprises a first interface, for data communication to and from data equipment at a plurality of customer premises. The switch also includes a second interface, for communication with a first network domain, such as the wide area internetwork, over a low-layer protocol link carrying subscribers' traffic in accord with a first type of a higher level protocol. A third interface enables communication with a second network domain, via a second type of the higher level protocol.

The inventive data switch also includes a switch fabric coupled to the interfaces and a controller coupled to the switch fabric, for controlling routing of traffic through the fabric between the interfaces. Upstream transmissions, received through the first interface, are segregated and routed either to the second interface or to the third interface, depending on whether the transmissions utilize the first or second type of the higher level protocol, respectively. Downstream communications, received via the second and third interfaces for respective customer premises data equipment, are aggregated and routed through the first interface. In the preferred embodiment, the data switch implements a logical circuit, for example as an ATM permanent virtual circuit provisioned through the switch fabric from the first interface to the second interface for each customer premises.

Hence, the preferred embodiments alleviate the above noted problems by providing an intermediate node, typically an enhanced switch, to segregate upstream traffic based on analysis of the type of communication. This analysis utilizes protocol information contained in each communication, for a protocol higher than the switching protocol, that is to say higher than a protocol used to

- define the customer's logical circuit. In an implementation utilizing ATM, for example, the logical circuit is a virtual circuit, and the node examines the protocol information for layers above the ATM layer. One type of traffic remains on the virtual circuit, whereas other traffic is handed off to a vertical services domain.
- 5     The node also provides a point to aggregate traffic from the from the vertical services domain with other traffic on the subscriber's logical circuit, for transport to the customer premises equipment.

To this end, the preferred embodiments use an ATM switch having routing/processing capabilities at the higher layer, for example at the layer-3  
10    protocol or above, in close proximity to a DLSAM. The links between the DSLAM and this switch can be of any desired high capacity because they are short and therefore relatively low in cost. Consequently, the DSLAM and the new ATM switch can support relatively high bandwidth from that node to the customer premises. To utilize the increased bandwidth and support new services, the  
15    modems on the opposite ends of the digital subscriber line negotiate and operate at the maximum rate that the line conditions will permit.

The ATM switch essentially subdivides the traffic associated with each subscriber line. One branch goes to a gateway router and hence to one or more ISP(s) at the rate corresponding to the Internet access subscription. It may be  
20    helpful to consider this as long distance or wide area traffic for the subscriber. The other branch is for local traffic, to and from the locally connected vertical services domain. The interconnection to the vertical services domain supports QoS and introduction of vertical services not easily provided from the public Internet, such as video on demand, multicasting, and voice over IP. The vertical  
25    services domain is relatively secure since it is not accessible from the public Internet.

In operation, the access network will still support a logical circuit in the form of a virtual circuit to the gateway router for each customer, for example in the form of an ATM permanent virtual circuit. That circuit is provisioned for the  
30    subscriber's desired grade or quality of Internet access service. The subscriber line, however, supports at least the same and often a higher rate service, for example via rate-adaptive ADSL communication over a twisted wire line. In the

- 18 -

upstream direction, the ATM switch examines at least one aspect of each data transmission relating to a logically higher level of protocol, e.g. in a layer 2 protocol encapsulated in ATM, or at layer 3, layer 4, or higher. From the higher level information, the switch determines the type of transmission and segregates  
5 the upstream transmissions on the subscriber's virtual circuit. One type of transmission continues on the virtual circuit, e.g. for Internet transport. Any and all other types, however, go to the vertical services domain. In the downstream direction, the ATM switch aggregates any communications coming from the vertical services domain together with downstream transmissions on the virtual  
10 circuit from the Internet.

The vertical services domain also represents a communication network. The vertical services domain, however, preferably takes the form of a data network optimized for local transport of vertically injected services, that is to say local data traffic. In view of its local nature, it is easier and more cost effective to  
15 provide high bandwidth services. The vertical services network, for example, could take the form of a giga-bit Ethernet type local area network. Also, it is easier to adapt the vertical services network to support service level agreements with customers with regard to quality of service. In many cases, it actually is sufficient to support QoS on the one hop through the ATM switch, itself.

20 To support the QoS requirements, a feature of the preferred embodiments involves certain queuing and tagging operations within the ATM switch. Essentially, the switch will maintain two or more queues for each permanent virtual circuit. The switch distinguishes the queues based on importance. As the switch receives cell transmissions for transport over the virtual circuit to the  
25 customer premises, the switch will internally tag each cell as to its importance level and place the cell in the appropriate queue. The switch may implement any one of a number of different algorithms to select and transmit cells from the various queues. The particular algorithm is selected to implement QoS in conformance with the subscriber's service level agreement with the carrier  
30 and/or agreements between the carrier and the vertical services providers.

Within the one virtual circuit assigned to the individual subscriber, the invention actually provides multiple tiers of service, preferably with multiple

levels of QoS. Also, at different sections along the virtual circuit "pipe," the network provides different levels of rate shaping. All layers and all services are available at the home, but different services receive different treatments in the network conforming to the different levels of QoS. The inventive approach, 5 however, does not require each subscriber to have multiple virtual circuits.

Services provided on the vertical services domain appear as IP data services. Virtually any communication service may utilize the vertical services network and through it to gain access to the carrier's local customer base, simply by providing an IP interface for coupling to the vertical services network. For 10 example, it is a simple matter to connect any digital source of broadcast audio or video information, such as a direct satellite broadcast receiver system similar to those used today in residential applications, through an IP interface. Such a broadcast source and interface can provide the full range of received video services, over the vertical services network. The access data network may 15 distribute the video programming to a number of access switches within a local geographic area. The switch provides an optimum point for frame or cell replication for multicasting services. Hence, in our video example, the switch replicates and distributes frames for the broadcast service over the digital subscriber line circuits to customers desiring to view the programming.

Another aspect of the present invention encompasses a method for 20 customer-centric selection of telecommunications services and a consequent method for automated provisioning of the customer's data network services. The preferred embodiment stresses a correlated information flow, in which the customer's media access control (MAC) address has a corresponding, temporary 25 IP address and a corresponding virtual circuit identifier. Specifically, a customer that is capable of receiving vertical services can use a web-based software application to order service. Software entities extant in a service domain coordinate the processes that will automatically provision the corresponding data circuit, according to both customer-selected parameters and detected network 30 parameters. This exemplary method involves the components described previously to enable the vertical services insertion within an ADSL data network, as well as components providing the relevant software resources.

The fundamental challenge addressed by such a provisioning methodology is to significantly reduce, or completely remove, the errors present in the existing data circuit provisioning process, which includes: manual ordering processes; manual record-keeping; and manual circuit changes at the PSTN frame and the intermediate ADSL frame.

A further aspect of invention relates to the unique software for implementing the automated service selection and provisioning procedure. A software product, in accord with this aspect, includes at least one machine readable medium and programming code, carried by the medium. In the preferred embodiment, the code actually includes several applications, such as a Services Application, a Metering Application and a Provisioning Application, which may reside in separate media and run on two or more servers or other network nodes, although the inventive concept encompasses operation from a single, common server.

A computer readable medium, as used herein, may be any physical element or carrier wave, which can bear instructions or code for performing a sequence of steps in a machine readable form. Examples of physical forms of such media include floppy disks, flexible disks, hard disks, magnetic tape, any other magnetic medium, a CD-ROM, any other optical medium, a RAM, a ROM, a PROM, an EPROM, a FLASH-EPROM, any other memory chip or cartridge, as well as media bearing the software in a scannable format. A carrier wave type of medium is any type of signal that may carry digital information representative of the instructions or code for performing the sequence of steps. Such a carrier wave may be received via a wireline or fiber-optic network, via a modem, or as a radio-frequency or infrared signal, or any other type of signal which a computer or the like may receive and decode.

Additional objects, advantages and novel features of the invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following and the accompanying drawings or may be learned by practice of the invention. The objects and advantages of the invention may be realized and attained by means of

the instrumentalities and combinations particularly pointed out in the appended claims.

Brief Description of the Drawings

5 The drawing figures depict preferred embodiments of the present invention by way of example, not by way of limitations. In the figures, like reference numerals refer to the same or similar elements.

Fig. 1 is a functional block diagram of a digital subscriber line data network supporting enhanced services in accord with the inventive concepts.

10 Fig. 2 is a slightly modified functional block diagram of network of Fig. 1, illustrating the protocol stacks used in the various network elements.

Fig. 3 is a functional block diagram of the network showing the service level agreements (SLAs) for which the network provides appropriate QoS.

Fig. 4 is a logical diagram of the functional elements of an L3/4 switch, for use in the inventive network of Figs. 1-3.

15 Fig. 5 is a block diagram of a modified portion of the network, useful in explaining migration to other types of physical transport and switching/routing protocols.

Fig. 6 is a block diagram of a portion of the network of Fig. 5, showing the interconnection thereof with the wide area network and the local vertical services 20 domain.

Fig. 7 is a block diagram of the network similar to those of Figs. 1 and 2, but showing details with regard to certain elements thereof enabling an automated service selection and provisioning procedure in accord with a further aspect of the invention.

25 Fig. 8 is a flow-chart illustrating the steps of an example of the automated service selection and provisioning procedure.

Fig. 9 a block diagram of a prior art asymmetrical digital subscriber line data network.

Fig. 10 is a slightly modified functional block diagram of the prior art network illustrating the protocol stacks used in the various network elements.

Detailed Description

The inventive network architecture introduces QoS into the ADN, in a manner that enables the delivery of sophisticated and demanding IP-based services to subscribers. The architecture does not affect existing Internet tiers of service such that the promised rates for such access appear the same as offered today. Also, the new architecture is cost-effective in terms of initial costs, build-out, and ongoing operations. The architecture utilizes an access switch or router capable of examining and selectively forwarding packets based on a relatively higher layer of the protocol stack, that is to say based on information encapsulated within the layer-2 information utilized as the lowest level definition of connectivity through the network. The access switch enables segregation of upstream traffic by communication type and downstream aggregation of wide area traffic together with traffic from a local vertical services domain.

In accord with the invention, the access switch examines the higher-level protocol information in the upstream communications to segregate traffic into two or more streams based on traffic type. A logical circuit provisioned for the subscriber extends from the customer premises through the switch to a node providing service to the packet switched network, typically on the Internet. The provisioning of this circuit defines a contiguous connectivity for the logical circuit, in the layer-2 protocol utilized by the network. Higher level protocols ride within this particular layer-2 protocol, including some protocols that are often considered themselves to be layer-2 protocols.

The analysis of upstream packets to determine higher-level type enables segregation of the upstream traffic into two or more upstream flows. For example, this allows the switch in the preferred embodiment to keep traffic bound for the packet switched internetwork on the logical circuit but route all other traffic onto another network, in this case a local network for vertical services. In the preferred embodiment, from the access switch or router to the

- 23 -

packet switched network, the provisioning for the logical circuit limits the rate(s) of communication to those defined by the grade of service to which the customer subscribes. From the customer premises to the switch, however, the communications run at the optimum rates that the facilities can physically support, e.g. the maximum rates that the ADSL modems or the like can adaptively establish for the particular subscriber's line.

In the downstream direction, the switch aggregates traffic for each subscriber. The switch receives the rate-limited traffic from the packet switched network, on the subscriber's logical circuit. The switch also receives any downstream traffic intended for the subscriber, from the vertical services network. The switch combines this traffic and sends the combined communications downstream over the subscriber's logical circuit to the customer premises, at the optimum downstream rate that the subscriber's facilities can support.

Fig. 1 provides a high-level functional illustration of an exemplary digital subscriber line network, specifically an ADSL data network 10, implementing the various concepts of the present invention. Fig. 2 provides an alternative functional illustration of the elements of such a network. Of particular note, Fig. 2 shows the various protocol stacks associated with the appropriate network elements that utilize those stacks. The end-user may be a single PC user or a small business or a residential LAN user. The data equipment of such users typically includes servers and PCs and may include a variety of other devices, such as fax machines, televisions, recorders and remote controlled appliances, having data communication capabilities.

The customer access link comprises an xDSL twisted pair, although those skilled in the art will recognize that the invention is readily adaptable to networks using other types of links to the subscriber premises. In the illustrated embodiment, the network 10 supports ADSL, which the carrier may offer in grades supporting 640 kbps, 1.6 Mbps or 7.1 Mbps (downstream) rates for Internet access. The actual communications over the DSL loops, however, run adaptively at the maximum rates that the line conditions allow.

The ADSL-based local access data network or "ADN" 10 provides access to two different network domains for communication services. The two network domains are logically separate. In most implementations, the first domain may be considered as a long distance or wide area domain, whereas the second 5 domain is a local network domain. In the illustrated example, the ADN 10 provides access to a first domain in the form of a wide area internetwork, such as the public Internet, corporate local area networks (LANs), and the like, represented by the network cloud 11 for the ISPs. The high speeds available through the local network 10 enable a wide range of communications, for 10 example, of text data, of video data, for multimedia, for web browsing, of transfers of files, for database searching, and the like via the network(s) 11.

In accord with the invention, the ADSL-based local access network 10 also offers access to a wide variety of other IP-based services through a local data network 13 serving as the vertical services domain (VSD). The vertical services 15 typically are high-end services requiring certain QoS levels and often having a certain local characteristic. Examples of the vertical services are discussed in more detail later.

The vertical services network offers an efficient domain from which the carrier can locally inject high-end services and/or services of other local 20 providers. Because the vertical services domain is separate from the public Internet, equipment providing the vertical services is subject to attacks directly from the public Internet.

As shown in Figs. 1 and 2, a central office (CO) 15 comprises one or more 25 DSLAMs 17 and L3/4 ATM switch 19. Elements of the CO 15 providing plain old telephone service (POTS) have been omitted for simplicity of illustration, since they are generally similar to those shown in Fig. 9.

The switch 19 is designated as an "L3/4" switch here as a matter of convenience, to illustrate that the switch 19 has the capability to make selective packet forwarding decisions based on protocol information at some level that is 30 above the L2 level that the switch normally utilizes to define or establish logical circuit connectivity. It will be recognized, however, that some of the protocols, although higher than the ATM type level 2 protocol used by the preferred switch

are themselves often thought of as level 2 protocols even though they are above or encapsulated in the ATM type level 2 information. Also, decisions as to the selective forwarding may be made in response to monitoring of any level of the protocol stack above the L2 level that the switch normally utilizes to define or establish logical circuit connectivity, for example from any level above ATM all the way up to the L7 application layer.

Returning to the discussion of Figs 1 and 2, for purposes of this discussion, assume that the equipment at the various customer premises connect directly to the CO 15 via twisted pair type copper wiring 21. In an actual implementation, 10 many customers may connect through such wiring to a remote terminal linked to the CO 15 via optical fiber. Other hardwired, optical or wireless implementations of the digital subscriber lines are discussed later. In the illustrated embodiment, each line 21 from a customer premises connects to an ATU-C within one of the DSLAMs 17.

15 On the customer premises side, the digital subscriber line circuit 21 connects to an ADSL terminal unit (remote) or ATU-R 23. The ATU-R 23 is a modulator/demodulator (modem) for communicating over a twisted wire pair type line 21, in accord with the ADSL protocols. The ATU-R in turn connects to customer premises equipment, shown by way of example as a PC 25 at each 20 location (Figs. 1 and 2). Those skilled in the art will recognize that the customer premises equipment 25 may include a wide range of other types of devices having data communications capabilities (see e.g., Fig. 3).

The ADSL user's normal telephone equipment (not shown) also connects to the line 21, either directly or through a frequency combiner/splitter, which 25 often is incorporated in the ATU-R. The normal telephone signals are split off at both ends of the line and processed in the normal manner.

For digital data communication purposes, the ATU-C and ATU-R modem units create at least two logical channels in the frequency spectrum above that used for the normal telephone traffic. One of these channels is a medium speed 30 duplex channel; the other is a high-speed downstream only channel. Such modems may implement either one of two techniques for dividing the usable bandwidth of the telephone line to provide these channels. One approach is to

- 26 -

divide the usable bandwidth of a twisted wire pair telephone line by frequency, that is to say by Frequency Division Multiplexing (FDM). The other approach uses Echo Cancellation. FDM uses one frequency band for upstream data and another frequency band for downstream data. The downstream path is then  
5 divided by time division multiplexing signals into one or more high-speed channels and one or more low speed channels. The upstream path also may be time-division multiplexed into corresponding low speed channels. With echo Cancellation, the upstream band and downstream band substantially over-lap. The modems separate the upstream and downstream signals by means of local  
10 echo cancellors, in a manner similar to that used in V.32 and V.34 modems.

The DSL modems may use a number of different modulation techniques to physically transport digital data streams. A number of implementations of the modems have used carrierless amplitude phase (CAP) modulation. Most current xDSL modems, however, utilize a discrete multi-tone (DMT) approach.

15        Returning to the discussion of the CO 11, the structure and operation of each DSLAM 17 is essentially the same as those of the DSLAM 111 in the embodiment of Fig. 9, except that the control functionality of the DSLAM 17 is somewhat different. The DSLAM 17 controls the ATU-Cs to implement a rate-adaptive ADSL service, to adapt operations so as to maximize data rates for the  
20 communications over the individual subscriber lines. Essentially, the ATU-Cs and ATU-Rs signal each other over the lines to synchronize their modes of operation at parameter settings, which achieve optimum data throughput. Also, the DSLAM 17 does not need to monitor or limit the line rates, but instead relies on the rate-adaptive control algorithm to maximize the rates achieved over the  
25 ADSL circuits or provide rate-shaping for the ATM virtual circuits. Other network elements limit rates, where necessary.

The L3/4 ATM switch 19 is co-located with the DSLAMs 17, within one central office 15. As a result, it is practical to connect the multiplexer within each of the DSLAMs 17 over a high-speed data link directly to an appropriate port of  
30 the ATM switch 19. Because these links are short, there is little or no cost imposed when implementing such links using wideband equipment. By itself, the co-location of the L3/4 ATM switch 19 with the DSLAM(s) 17 does not increase

bandwidth. Rather, it makes increased bandwidth at the network edge economically feasible, due to proximity. Co-location removes the requirement to purchase expensive wide area transport (the SONET ring) to increase bandwidth. In particular, the direct OC3/OC12 connections between the DSLAM 17 and the 5 L3/4 ATM switch 19 do not incur any recurring tariff expenses.

The ATM switch 19 connects through a SONET ring 27 to a gateway router 29 providing ATM transport through a cell relay network 30 (Fig. 2) to the ISPs shown at network 11 in the drawings. Most of the ISPs will utilize a concentrator or other equipment as their point of presence for Internet access (Fig. 2). In the 10 preferred embodiment, the equipment 31 provides a point-to-point protocol (PPP) interface designed for transport over Ethernet (PPPoE). The ATM switch 19 also provides a connection to the local implementation of the VSD network 13, for example via a giga-bit Ethernet port to a switch or other local network elements 18.

15 The illustrated local access type ADN network 10 provides ATM cell transport from the customer premises to the ISP network(s) 11. The ATM cells serve as the layer-2 protocol for defining contiguous switched connectivity. Higher level routing protocols, such as Ethernet and TCP/IP frames, ride within the ATM cells. Services of different types utilize different protocols at one or 20 more layers above the ATM cell layer. In the preferred embodiments, all communications utilize Ethernet. However, communications to and from the ISPs use the noted PPPoE type Ethernet protocol. In contrast, communications to and from the vertical services domain use one or more of the other Ethertype protocols.

25 To efficiently provide cell relay, each customer is assigned a virtual circuit that extends from the ATU-R 23 in the respective customer premises to the gateway router 29. This logical circuit is defined at the layer-2 protocol level. The present implementations implement this logical communication circuit as an ATM permanent virtual circuit, although the inventive concepts may apply to 30 other types of logical circuits or channels.

The gateway router 29 is the communication node of the access network 10 providing access to the wide area IP packet networks, of corporations or more

often of Internet Service providers. The gateway router 29 terminates permanent virtual circuits through the cell relay network 30, from the equipment 31 of each such wide area packet network provider 11. The gateway router 29 also terminates the permanent virtual circuits from the subscribers through the data network 10. For communication with a selected ISP network 11, for example, the gateway router 29 routes cells from the permanent virtual circuit from the subscriber through to the permanent virtual circuit of the selected ISP network 11. In the opposite direction, the gateway router 29 routes cells from the permanent virtual circuit from the selected ISP network 11 through to the permanent virtual circuit of the particular subscriber.

For the leg of the subscriber's logical circuit, extending from the L3/4 ATM switch 19 through the gateway router 29, the carrier programs one or more nodes along the path behind the DSLAMs 17, to regulate traffic on the virtual circuit to the rate corresponding to the grade of Internet access service to which the particular customer subscribes. In the preferred embodiment, at least one such node performing this rate shaping function is the L3/4 ATM switch 19. All traffic going to and from the ISP network(s) 11 therefore is still limited to the rates defined in the service level agreement (SLA) for Internet access that the carrier has with the particular customer.

The portion of the virtual circuit extending between the ATM switch 19 and the ATU-R 23, however, is not rate limited but instead runs at the maximum rate that the line will support using the rate-adaptive ADSL modem operation. In most cases, the rate-adaptive ADSL modem operation will support rates substantially higher than the subscription rate for Internet access service.

The L3/4 ATM switch 19 also provides the interconnection to the subscriber's virtual circuit for insertion of downstream traffic from the vertical services domain 13 and separation of upstream traffic from the subscriber's virtual circuit going to the vertical services domain 13. In accord with the invention, decisions as to whether upstream traffic is destined for the vertical services domain 13 or should remain on the subscriber's virtual circuit going through the gateway router 29 and the cell relay network 30 to the ISPs 11 is based on an analysis of traffic type. The traffic type analysis relies on protocol

information contained in the communications, which relates to layers of the protocol stack that are higher than the layer-2 switching protocol, in this case above the ATM layer.

As shown in Fig. 2, traffic destined for an ISP 11 utilizes a variation of a point to point protocol (PPP) intended to run on top of Ethernet, referred to as PPP over Ethernet or "PPPoE." A 'type' indicator contained within the Ethernet frames identifies the PPPoE protocol. In contrast, traffic going to and from the vertical services domain utilizes other 'types' of Ethernet protocol. All traffic to and from the customer premises uses Ethernet frames carried within ATM cells.

The switch 19 therefore routes a subscriber's traffic going to and from an ISP 11, upon detection of the PPPoE indicator in the level 3 data contained within the Ethernet cells. This traffic will also utilize public IP addressing. In contrast, the ATM switch 19 routes a subscriber's traffic going to and from the vertical services domain, upon detection of any other type of Ethernet protocol at level 3 or above in the protocol stack. The IP addressing in the vertical services domain 13 utilizes private-IP addresses, for example, as administered with a DHCP server (not shown) coupled to the network cloud 33. Although shown separately, the cloud 33 may be implemented as a portion of the network providing the physical elements of the vertical services domain. The portion 33, however, would typically be a logically separate domain that the carrier controls and restricts for its own network administration use.

Fig. 3 depicts the logical division of the subscriber's traffic, as implemented at the ATM switch 19 in accord with the invention. As shown, the network 10 provides a logical "pipe" or circuit 35 extending to the networks 11 of one or more of the ISPs, for an Internet access application. The ATM switch 19 (Fig. 1) preferably performs a rate shaping or control function. The leg 35 of the subscriber's traffic extending to the ISP 11 provides upstream and downstream communication rates conforming to a service level agreement (SLA) applicable to the subscriber's Internet access application. As such, the communications over the subscriber's logical circuit, extending from the switch to the ISP, provide a first level of QoS. To the subscriber, service over the leg 35 appears the same as a subscriber selected grade of Internet access service as provided by older ADN

- 30 -

architectures. Fig. 3 illustrates chat rooms, web surfing and e-mail as examples of services an ISP might offer through the Internet Application SLA circuit 35 and the attendant portion of the subscriber's assigned logical circuit through the access network.

5       The network 10 also supports communications over one or more logical application paths 36 to local applications 37 hosted in the vertical services domain. Assuming that a subscriber with various equipment 25 also subscribes or otherwise participates in one or more of the vertical services, the local carrier (e.g. Bell Atlantic in Fig. 3) offers a corresponding number of additional  
10 application SLAs with the customer. Each SLA for a vertical service may specify QoS parameters for the particular application, such as rate/bandwidth, latency, jitter, packet loss, packet sequence, security and/or availability. Examples of such applications hosted in the carrier's vertical services domain 37 include the illustrated voice over IP service shown as a V/IP gateway, as well as video services  
15 and some caching for high volume local web services. Communications for such applications utilize the one or more paths 36.

The network of Figs. 1 to 3 also supports closed or private user work groups or virtual local area networks (VLANs). VLANs, for example, may be implemented as a service hosted through the vertical services network 13. All  
20 traffic from customer premises belonging to one of the VLANs would include a VLAN Ethertype indicator. Since the frames would not be PPPoE type Ethernet frames, the switch 19 would route the upstream transmission to the vertical services network 13. VLAN frames also carry a VLAN identifier. The nodes of the network 13 transport the packets for the VLAN members only to resources that  
25 are members of the identified VLAN. In this manner, the capability exists for creating logical workgroups of users and their shared resources (servers, printers, etc.), that may be physically separated from each other. VLAN groupings can provide privacy and security to their members while enabling "broadcast domains" whereby broadcast traffic is kept logically "inside" the VLAN.

30       The present invention also supports segregation and aggregation of traffic for three or more domains, based on the higher-level traffic type analysis. An alternative implementation for closed user group services therefore might

actually provide a third domain for such services. PPPoE traffic would go to the concentrator equipment 31, as discussed above. The switch 19 would route the upstream transmission frames bearing a VLAN Ethertype indicator to the alternate network serving as the VLAN service domain. That domain would 5 process VLAN communications essentially as done in a normal switched Ethernet type network. Upstream transmissions that bear any other type of Ethertype indicator would go to the vertical services network 13, as in the earlier examples.

A feature of the switch, in accord with the invention, is that it prioritizes traffic for each customer to support QoS for the various services as required by 10 service level agreements (SLAs) between the customer and the carrier. In this regard, one implementation of the L3/4 ATM switch 19 performs queuing and tagging operations in accord with the desired prioritization. The switch will maintain two or more queues for each subscriber's permanent virtual circuit. The switch distinguishes the queues based on importance or priority. As the switch 15 receives cell transmissions for transport over the virtual circuit to the customer premises, the switch will internally tag each cell as to its importance level and place each cell in the appropriate queue based on the tag.

The tagging and prioritization may be based on traffic type or 'Type of Service' (ToS). Table 1 illustrates one example of the possible ToS levels that may 20 be assigned to different communications.

- 32 -

Relative Priority	ToS Value	AR Queuing	Customer Traffic <i>Encapsulated TOS</i>	Internal BA Mgmt Traffic <i>Rewritten TOS</i>
Critical	0	WFQ		OSPF, SNMP, ICMP,BGP
Management		(Control 25%)		
Future	1			
Real Time Interactive	2	WFQ (High 40%)	Port numbers identified	
IP Application Control	3	WFQ (Medium 30%)	ICMP, IGMP, EGP, DNS, H.323 signaling, BGP, SIP, Microsoft Media Player Streaming Media Control, RTSP	IGMP, RADIUS
One Way Streaming Media	4		UDP (ports 1024+)	
One Way Batch	5		HTTP, HTTPS, SNMP, Telnet	
Unknown	6	WFQ	<i>Other</i>	<i>Other</i>
Non time sensitive	7	(Low 5%)	FTP, TFTP, SMTP	

Table 1

The access switch 19 will examine the content of each communication and determine an appropriate ToS level, for example in accord with the table above.

- 5 Based on the ToS level, the switch will add a tag to the cell(s) as part of its

internal processing. Using the ToS tags, the switch will place each of the cells for a given subscriber into a corresponding one of a plurality of queues that the switch maintains for the subscriber's traffic.

The switch may implement any one of a number of different queue servicing algorithms to select and transmit cells from the various queues. For example, the switch 19 may implement one of several statistical algorithms, equal queuing, weighted queuing, priority selection from queues, etc. The particular algorithm is selected to implement QoS in conformance with the subscriber's service level agreements (SLAs) with the carrier. In most cases, particularly for service applications from the vertical services domain, the switch 19 will not normally drop any cells or packets. In the rare event that the switch 19 becomes congested, any dropping of cells is based on the priority level assigned to the frame, i.e., lowest priority first. Also, if the switch ever drops cells, it drops all cells for an effected lowest priority frame from the relevant queue.

The ability to prioritize traffic across the inserted streams and the PPPoE stream enables the operator to control flows through the ADN 10 so that the local access facility is not overwhelmed with content which exceeds its physical (rate adaptive) limitations. For example, the queuing rules preferably ensure that the 'proper' applications (based on insertion device based rules) obtain access to the limited rate adaptive bandwidth available on any given subscriber's line. Also, the insertion point, switch 19, will usually sit behind the media conversion point (e.g., the DSLAM 17). An OC3 or other facility between the switch 19 and the DSLAM 17 also could become congested. Preferably, the switch 19 or other element at the insertion point queues the traffic in such a manner that no downstream facility (OC3) limitations (which are independent of the rate adaptive DSL loop limitations) will result in packets being dropped.

The queuing will be done based on customer and network provider determined rules so that contention for the facilities facing the subscriber will be addressed via the dropping of the appropriate packets. That way an inserted video stream doesn't overwhelm a PPPoE or Voice stream (due to facility limitations). Among others, appropriate industry understood methods for

accomplishing this queuing control include Weighted Fair Queuing (WFQ), Priority (PQ) Queuing, and Weighted Random Early Discard (WRED).

Also, the ability of the switch 19 to examine higher level information provides other advantages in network operations. For example, the switch can 5 implement sophisticated filters on the higher level information, e.g., to provide security. As another example, the switch preferably performs measuring and monitoring to determine what if any packets get dropped (based on the physical rate adaptive limitations), and generates appropriate reports to an external operations system (not shown).

10        The introduction of the L3/4 ATM switch 19 in proximity to the DSLAM(s) 17 also provides benefits in terms of operation of the gateway router 29. Due to the end-to-end use of the Unspecified Bit Rate (UBR) PVCs, the gateway router interface to the cell relay network 30 has been engineered to support a maximum 15 of 2000-4000 PVCs (end users). This is essentially an over-provisioning of bandwidth that probabilistically avoids service degradation that could result from simultaneous demand for bandwidth. The ability of the L3/4 ATM switch 19 to perform QoS and rate shaping essentially reduces or ever removes this concern, because it significantly reduces the risk that the gateway router 29 will become a bottleneck. As a result, the ADN 10 can increase bandwidth efficiencies for this 20 interface. Further, the capacity through the gateway router 29 need not be upgraded as often to support demand for increased bandwidth associated with new bandwidth-intensive services, since many such services are now introduced through the vertical services domain 13 and the L3/4 ATM switch 19.

25        To fully understand an exemplary implementation of the various inventive concepts, it may be helpful to consider an ATM-based embodiment of the L3/4 switch 19. Fig. 4 is a block diagram of the elements and functions of such a preferred embodiment of the switch 19.

30        The preferred embodiments utilize Ethernet framing. As shown in the drawing, the switch 19 includes an Ethernet interface 41, an ATM interface 42 and an associated physical interface 43 facing toward the subscribers. In an embodiment for use in the network of Figs. 1 and 2, the physical interface might take the form of one or more OC-3 or OC-12 links to the DSLAMs 17. These links

carry all ATM cell traffic going to and from the DSLAMs and hence to and from the customer equipment served through the particular switch 19.

The switch 19 also includes an Ethernet interface 44, an ATM interface 45 and associated physical interface 46 facing toward the gateway router 29 and hence the ISPs 11. The physical interface 46 might take the form of one or more OC-12 or OC-48 links to the gateway router 29. These links carry all ATM cell traffic going to and from the ISPs or other wide area inter-networks 11. For these communications, the Ethernet interface 44 passes through PPPoE traffic, as specified by the Ethertype indicator in the cells transporting the relevant frame segments.

Facing the vertical services domain, the switch 19 includes an Ethernet interface 47 and a physical interface 48. These interfaces conform to the particular network utilized by the carrier for the vertical services domain, such as giga-bit Ethernet over wire or optical links.

The switch fabric 49 performs the physical switching of data along various paths through the switch 19, in response to instructions from a programmed routing controller 50. Fig. 4 also shows the communications flow through the switch, for each subscriber. The switch 19 also implements a Decision Point 51, shown for example within the Ethernet interface processing 41 on the subscriber facing side. At that point, the PPPoE traffic is separated from all other traffic. From that point, the PPPoE Flow 52 for each subscriber extends as a portion of the subscriber's ATM virtual circuit, facing the cell relay network and hence the ISPs 11. The PPPoE Flow 52 contains Ethernet frames that are of PPPoE Ethertype. Facing towards the subscriber premises, the switch 19 implements an Aggregate Flow path 53, in the form of another portion of the ATM virtual circuit, which contains all ingress/egress subscriber traffic. The switch implements a Generic Path 54 extending through the interfaces to the vertical services network. This path 54 carries all traffic other than PPPoE.

In this example, the switch 19 implements the Decision Point 51 based on recognition of the Ethertype indicator, which is above the layer-2 ATM cell routing information. However, the switch may implement the Decision Point 51 based on still higher-level protocol information.

- 36 -

In a preferred embodiment, the Ethernet and ATM interfaces 41 and 42 and the Ethernet and ATM interfaces 44 and 45 implement segmentation and reassemble (SAR) functions, essentially providing two-way conversions between ATM cell format and Ethernet frame format. Segmentation involves dividing an  
5 Ethernet frame into a number of 48-byte blocks and adding ATM headers to the blocks to form a corresponding number of ATM cells. Any blocks that do not include a complete 48-byte payload are padded as necessary. Reassembly entails receiving and buffering ATM cells until it is recognized that a complete frame has been received. The ATM headers of the cells and any padding are stripped, and  
10 the payload data is reassembled into the form of an Ethernet frame.

In such an embodiment of the switch 19, the decision point 51 determines how to selectively forward the Ethernet frame information taken from a particular series of upstream ATM cells based on the Ethernet information taken from the ATM cell payloads, for example, by examining the frame header and  
15 recognizing the particular Ethertype indicator. Internally, the actual switch fabric 49 for such an embodiment of the switch 19 would comprise an Ethernet switch, even though to other elements of the ADN network 10 the switch 19 appears to perform an ATM switching function.

Those skilled in the art will recognize however, that the decision and  
20 switch fabric may be implemented in other ways. For example, a series of cells corresponding to an Ethernet frame could be buffered and the payloads examined just to recognize and identify the Ethertype indicator, without a complete reassemble of the Ethernet frame. This later implementation therefore could utilize an ATM cell-based switch fabric.

25 From the discussion above, it should already be apparent that certain aspects of the invention relate to setting up logical communication circuits at a relatively low protocol layer corresponding to switching or routing functions and then segregating traffic by distinguishing communication type using higher level protocol information. To insure full understanding on these points, it may be  
30 helpful to consider the protocol layer definitions, with particular reference to the illustration of the preferred layers in Fig. 2. The International Standards Organization (ISO) Open Systems Interconnection (OSI) reference model

specifies a hierarchy of protocol layers and defines the function of each layer in the network.

The lowest layer defined by the OSI model is the physical layer (L1). This layer provides transmission of raw data bits over the physical communication channel through the particular network. For example, on the subscriber lines in the preferred embodiment, the physical layer (L1) uses ADSL. Within the customer premises, communications use an Ethernet physical layer (L1), such as 10Base-T. Upstream network elements may use DS3 at some points, but most use SONET, for example OC-3 or OC-12 physical layer transport. Attention is directed to the lower half of the diagram in Fig. 2, which depicts the various protocol stacks throughout the network 10.

The layer defined by the OSI model next to the physical layer is the data link layer (L2). The data link layer transforms the physical layer, which interfaces directly with the channel medium, into a communication link that appears error-free to the next layer above, known as the network layer (L3). The data link layer performs such functions as structuring data into packets or frames, and attaching control information to the packets or frames, such as checksums for error detection, and packet numbers. In the network 10, the data link layer (L2) is used to define certain switching functions through the network. The network layer (L3) provides capabilities required to control connections between end systems through the network, e.g., set-up and tear-down of connections.

The preferred embodiments utilize ATM cell transport as the lowest element of the data link layer (L2), for example to define the connectivity extending from the ATU-Rs 23 through the ADN network 10 to the ISP or corporate networks 11. Subscriber virtual circuits are provisioned at the ATM cell layer, that is to say at the data link layer (L2). Similarly ISP virtual circuits are provisioned at this ATM data link layer (L2), from the gateway router 29 through the cell relay network 30 to the ISP access concentrators 31. The ATM protocol therefore is the layer-2 (L2) protocol used to define the logical connectivity from the subscriber premises to the gateway router 29. The ATM protocol also is the layer-2 (L2) protocol used to define the logical connectivity from the gateway router 29 to the ISP concentrators 31.

For purposes of this discussion, higher level protocols are protocols that ride on or are encapsulated within the particular layer-2 protocol, that is to say in the payloads of the ATM cells in the preferred embodiment. Such higher level protocols include some protocols, which are often considered themselves to be  
5 level-2 protocols, where they are transported within ATM cells. The preferred embodiments use Ethernet, a local area network protocol above the ATM portion of the L2 layer. Technically, the Ethernet protocol may be considered as another L2 layer protocol. However, because it is segmented and encapsulated into the payloads of the ATM cells, the Ethernet protocol information actually is a higher  
10 level protocol information above the specific level-2 protocol (ATM) that defines the normal connectivity through the ADN network 10.

In the OSI model, a transport layer protocol (L4) runs above the network layer. The transport layer provides control of data transfer between end systems. Above the transport layer, a session layer (L5) is responsible for establishing and  
15 managing communication between presentation entities. For example, the session layer determines which entity communicates at a given time and establishes any necessary synchronization between the entities. Above the session layer, a presentation layer (L6) serves to represent information transferred between applications in a manner that preserves its meaning  
20 (semantics) while resolving differences in the actual representation (syntax). A protocol (L7) that is specific to the actual application that utilizes the information communicated runs at the top of the protocol stack.

In accord with one inventive concept, the network 10 actually utilizes two or more different types of protocol at levels above the protocol within the L2 layer  
25 that actually defines the network connectivity. The ADN network 10 may use different protocols at the higher layers as well. By distinguishing transmissions based on differences in these higher-level protocol types, the ATM switch 19 separately forwards different types of communication traffic for each subscriber. In the preferred embodiment, communications to and from the ISP or corporate  
30 networks 11 utilize point-to-point protocol (PPP) as the network layer (L3) protocol and a shim for transport of PPP over Ethernet (PPPoE). PPPoE, as one Ehtertype protocol could also be considered as a second layer (L2) protocol albeit

above the Ethernet layer itself, which in turn rides on the ATM cells used for routing at least through the permanent virtual circuit at the L2 layer.

In the illustrated implementation, however, the use of the PPPoE or a different protocol actually is an indication of a difference in type of the higher layer protocols. In the illustrated example of Fig. 2, the vertical services domain traffic utilizes Ethernet (802.3 SNAP) above the ATM adaptation layer (AAL). As noted, the presently preferred L3/4 switch 19 implements its routing decision based on recognition of the Ethertype indicator, that is to say to distinguish the PPPoE traffic from all other types of transmission from the customers' data equipment. In view of the use of ATM as the data link layer (L2) protocol of the network defining the lowest layer of network connectivity for communications services through the ADN network 10, the discrimination based on Ethernet actually implements a decision based on an effectively higher protocol layer.

IP protocol carries the actual higher-level applications information, for transport to and from the vertical services domain and for transport to and from the wide area internetwork. As such, IP and its related transport protocol referred to as the "Transmission Control Protocol" (TCP) ride on top of (are actually encapsulated within) the lower level protocol elements discussed above. Presentation and application layer elements ride on top of the IP layer. IP communication requires that each user device have an assigned IP address. IP addresses, however, are a scarce commodity. Because of the use of IP transport for both wide area services and vertical domain services, the network 10 actually may at times assign two different IP addresses to each active data communication device of an end-user, albeit on a temporary basis. The wide area communications and the vertical services network may also be viewed as two separate 'broadcast' domains.

First, the carrier operating the ADSL data network 10 and the vertical services domain network 13 will maintain a pool of local addresses for assignment, on an as-needed basis, to end user equipment 25. To the carrier, the available IP addresses are a limited resource. Accordingly, the carrier assigns IP addresses on a dynamic basis, only to those users actually on-line at any given time. The carrier preferably utilizes private network type IP addresses and

dynamically administers such addresses using dynamic host configuration protocol (DHCP). DHCP is a protocol for automatic TCP/IP configuration, which enables dynamic address allocation and management.

When a particular device 25 becomes active via the ATU-R 23 and the  
5 DSLAM 17, it will activate a basic protocol stack, including an IP portion enabling communication with a DHCP server. The device will transmit an address request upstream through the network on the subscriber's virtual circuit. At the Ethernet level, this transmission appears as a broadcast message. The L3/4 ATM switch 19, however, will recognize that the packet is not a PPPoE communication and  
10 route the cells carrying the packet into the vertical services domain 13. Typically, a DHCP server is coupled to the vertical services domain network 13, for example as part of the carrier's administrative network or systems 33. The DHCP server somewhere on the vertical services domain 13, 33 will answer that broadcast request by selecting and providing an available one of the private IP addresses  
15 from the carrier's pool of available addresses. The message with the assigned address will go back to the L3/4 ATM switch 19 for insertion into the virtual circuit and transport back to the requesting device 25.

The particular end-user's device 25 uses the assigned private IP address as its source address, for all of its communications with the vertical services network  
20 13, so long as it remains on-line for the present session. When the overall session ends and the end-user device 25 goes completely off-line, the DHCP server returns the private IP address to its pool of available addresses, for reassignment to another user as the next user comes on-line.

As noted, the user equipment 25 receives a private IP address from the  
25 DHCP server. The addresses of services on the vertical services domain also are private IP networks. Because these addresses are private, they are accessible only to equipment within that domain and the data network 10. Consequently, the devices are not accessible to hackers or the like coming in through the public Internet.

30 This dynamic assignment of IP addresses allows the carrier to limit the number of IP addresses used to the number of users actively connected through the ISP's host to the Internet. The use of private IP addresses allows the user

equipment to communicate with the vertical services domain utilizing a normal IP-Ethernet protocol stack.

For the as-desired Internet access service, for example using a PPP or similar protocol, IP addresses are administered through the ISPs. The PPPoE 5 protocol preserves or emulates the traditional dial-up approach to ISP access. However, the PPPoE approach does utilize Ethernet and follows Ethernet standards, for example, involving processing of certain broadcast messages.

The user can select an ISP of choice, and her data equipment 25 will initiate a selective session through the Ethernet layer on the network 10 to access 10 the equipment 31 of the selected ISP network 11, in a manner directly analogous to a dial-up modem call through an ordinary telephone network. Hence at a time after initial activation through the networks 10 and 13, the user may activate a browser or other program for using the wide area internetwork service. This activates a second protocol stack, which includes the PPP protocol and the PPPoE 15 shim. The user selects an ISP, and the data equipment initiates communication through the network 10 to the PPPoE equipment 31 of that ISP.

The IP addresses used by each ISP are public network type IP addresses. To the ISP, the pool of available public IP addresses also is a limited resource. Accordingly, each ISP prefers to assign IP addresses on a dynamic basis, only to 20 those users actually on-line at any given time. Typically, as part of each initial access operation for a PPPoE session, the user's equipment 25 and the PPP terminating equipment 31 of the ISP conduct a handshaking, to establish data communications therebetween. As part of this operation, the user's device transmits a broadcast request for a public IP network. The broadcast message, in 25 PPPoE goes through the virtual circuit to the gateway router 29 and through the router and cell relay network 30 to the ISPs PPPoE equipment 31. Although it is a broadcast message, the network effectively limits transport thereof to the virtual circuit going to the ISPs PPPoE equipment 31, that is to a domain separate from the vertical services network domain 13.

30 The ISP host equipment 31 initiates a procedure to assign the user's computer 25 a numeric Internet Protocol (IP) address from the pool of available public addresses and sends a PPPoE message containing that address back to the

subscriber's device 25. When the session ends and the user goes off-line, the ISP host can reassign the address to another user, as the next user comes on-line.

This dynamic assignment of IP addresses allows the ISP to limit the number of public IP addresses used to the number of users actively connected 5 through the ISP's host to the Internet. The end-user equipment will implement a second protocol stack, carrying PPPoE communications. The PPP protocol will allow the end-user equipment to obtain and utilize the public IP address for communications going to and from the public internetwork.

The switch 19 will limit transport of other types of PPPoE broadcast 10 messages to the link to the PPPoE concentrator 31, in a manner similar to that described above for the PPPoE address request. The switch 19 also limits transport of non-PPPoE broadcast messages to the vertical services domain network 131, both for the address request message and for other types of broadcast requests. As such, the logical circuit to the PPPoE concentrator 31 15 becomes the conduit to one broadcast domain for upstream PPPoE messages; and the vertical services network 13 defines a second broadcast domain for upstream messages of other Ethertypes.

As noted, the end-user equipment 25 will implement two protocol stacks, a native stack without PPPoE and a second stack with PPPoE and a shim. In actual 20 operation, both the native stack with other Ethernet protocols and the wide area stack with PPP and the PPPoE shim often will be active at the same time. The software in the data equipment 25 will utilize one stack or the other depending on whether the user selected a link, e.g. a URL, relating to the wide area internetwork or the vertical services domain. For example, a browser may 25 display a page with embedded links. If a link is to a service on the vertical services domain, the embedded address will be a private address on the vertical services domain. Selection of such a link causes the device 25 to use the native Ethernet stack (without PPP or PPPoE) and the private address. Hence the L3/4 ATM switch 19 routes the request triggered by selection of the link to the vertical 30 services domain 13. In contrast, if the link is to a service on the public Internet or other network 11, the embedded address will be a public IP address. Selection of such a link causes the end-user device 25 to use the PPP and PPPoE stack and the

public address. Hence the L3/4 ATM switch 19 routes the request triggered by selection of the link over the virtual circuits to the PPPoE equipment 31 of the currently selected access provider network 11.

Services provided on the vertical services domain therefore appear as simple IP data services, albeit using the appropriate address space. Virtually any communication service provider may access the vertical services network 13 and through it the carrier's local customer base simply by providing an IP interface for coupling appropriate equipment to the vertical services network. For example, it is a simple matter to connect a direct satellite broadcast receiver system, similar to those used today in residential applications, through an IP interface to provide the full range of received video services over the vertical services network 13. This network may distribute the video programming to a number of L3/4 ATM switches 19, within a local geographic area serviced by the ADN network 10. The L3/4 ATM switch 19 co-located with the DSLAM 17 provides an optimum point for frame or cell replication for multicasting services.

For a multicast service, such as the satellite-originated video broadcast service, the service provider sends one stream through the vertical services domain network 13 to the L3/4 ATM switch 19. The switch 19 will monitor every ATM virtual circuit going to the subscribers, looking for IGNP requests. A subscriber sends an IGNP request to join a selected multicast channel. When the L3/4 ATM switch 19 detects such a request, it identifies the requested channel and the requesting subscriber equipment and forwards a 'join' message to the vertical services domain. Subsequently, the switch 19 replicates received packets for the requested broadcast channel, and the switch drops the replicated packets into the cells for each of the virtual circuits of all of the joined subscribers, including the newly added subscriber. When the subscriber later elects to end viewing of the multicast, the subscriber's equipment sends a 'leave' message, and the switch 19 stops adding the cells for the multicast to that subscriber's virtual circuit.

In addition to vertical services, the carrier continues to provide agreed access services to the equipment of the ISPs, in a manner analogous to current practices. For example, the carrier may provide its Internet access service to a

subscriber on a monthly subscription basis, at one of several available rates corresponding to the grade of internet access service (and thus the rate of communication to/from the ISP) selected by the customer's subscription.

In an enhanced service offering, the broadcast provider could offer a convenient navigation interface from a web server. The server could be on the vertical services network, but preferably is on the wide area Internet 11. With a PPPoE session active, the user can surf to the provider's server and view information about available programming. The user might select a current broadcast program by 'clicking' on a URL link in the provider's web-based information. Although provided through the wide area Internet 11, the URL would actually contain the private IP address for the desired broadcast program available from the vertical services network 13. Selection of such a URL therefore would generate a message to the appropriate server on the vertical services network 11 to initiate the above discussed procedure to allow the user to 'join' the selected broadcast. A similar methodology might also enable a provider to offer menu, selection and order/billing services from the Internet 11, to provide pay-per-view or video on-demand type services from the vertical services domain network 13.

Although IP-based, the services from the vertical services domain 13 may follow any other desirable business model. For example, a multicast service provider may contract with the carrier to provide multicast audio (radio-like) and/or video (TV-like) services via the vertical services domain. The multicast service provider, not the subscribers, would pay the carrier. The multicast service provider may offer any or all of the multicast programming to customers on some type pay-per-view basis but would likely offer most of the programming service for free or bundled in as part of some nominal monthly subscription charge. The multicast service provider instead would charge advertisers in a manner analogous to current broadcast business practices. Advertising distributed with the IP multicasting, however, can be carefully targeted at end-customers having demographic profiles meeting specific criteria specified by individual advertisers, which allows the multicast service provider to charge premium advertising rates.

- 45 -

Table 2 summarizes the characteristics and requirements of a number of examples of the types of vertical services that the VSD network 13 can deliver via the L3/4 ATM switch 19 and the ADSL data network 10.

<b>Vertical Service Offering</b>	<b>Characteristics of the Service</b>	<b>Network Requirements</b>
Voice Services	Local Co-Located VoIP Gateways, VoIP, Unified messaging, IP PBX, IP Centrex	Low latency, low jitter, non-correlated packet loss, and high availability
Video On Demand (Unicast)	Local VOD Servers or access to centralized servers.  Supports whatever model of server deployment/content delivery mechanism.	High bandwidth, low jitter, high availability, and low packet loss
Multimedia Broadcast (Multicast)	Broadcast Video; Broadcast Audio; Satellite Down Link support; Local Servers at the edge.	Varies with content type and with multicast implementation
Caching Services	Local servers at the insertion point, Local delivery mechanism for generic media objects such as web pages, images, video files, audio clips, software downloads, etc.	Layer 3/4 visibility
Distance Learning (EVC)	Integrated interactive video, voice and data	Low latency, low jitter, non-correlated packet loss, and high availability
Telecommuting	Closed user group with access to Transparent LAN Service (TLS).	IEEE 802.1Q

Table 2

The above discussed preferred embodiments implemented the processing above the layer-2 protocol in an enhanced ATM switch and focused on 5 implementation over an xDSL network specifically designed for use of twisted pair wiring to the customer premises. Those skilled in the art, however, will

- 47 -

recognize that the principles of the present invention are equally applicable to other types of layer-1 and layer-2 transport/switching technologies as well as selection based on other protocols above the layer-2 protocol. Fig. 5, illustrates the implementation of the layer 3/4 and higher switch functionality in a generic access router (AR) 61. The illustration also teaches the provision of digital subscriber line data communication between the access router (AR) 61 and a number of customer premises, using a variety of line technologies. The digital line technologies include dial-up modems 63, 65 as well as wireless communications between wireless asymmetrical subscriber loop (WASL) transceivers 67, 69. The access router (AR) 61 can service residential customers via these other communication technologies as well as through the DSLAM 17 and the ATU-R 23 as in the earlier embodiment. The access router (AR) 61 also serves business customer router equipment 71, using a variety of fast frame/cell packet technologies 73-76 and even optical fiber (SONET) 71.

Those skilled in the art will recognize that even these examples are limited. For example, the invention may apply to pure video networks, such as in a hybrid fiber-coax implementation of a CATV system with digital video service as well as cable modem service.

The access router (AR) 61 will provide one or more types of logical circuits, implemented in the appropriate layer-2 protocol(s), e.g. ATM, frame relay, etc. Although the links to the wide area internetwork and the vertical services domain have been omitted here for simplicity of illustration, the access router (AR) 61 will provide the routing functions to and from the wide area internetwork and the vertical services domain in a manner similar to the functionality of the L3/4 ATM switch 19 in the earlier embodiment. In this regard, the access router (AR) 61 will support the QoS levels and will enable local insertion of vertical services.

Fig. 6 depicts a portion of the network of Fig. 5, showing the interconnection thereof with the wide area internetwork and the local vertical services domain. The vertical services network itself may include a number of routers (R) 73. Through that network, the access router (AR) 61 provides communications with services in the VSD that may be purely local, somewhat distributed or even centralized. True long distance data services, such as chat

- 48 -

rooms, email and web browsing on the public Internet, however, are considered as Off-Net services, since they are accessed via the Internet access connection under the associated SLA.

Another feature of the present invention relates to a technique for  
5 customer-centric selection of telecommunications services and associated automated provisioning of the customer's data network services. The preferred embodiment of this feature of the invention utilizes a correlated information flow, in which the customer's MAC address has a corresponding, temporary IP address and a corresponding virtual circuit identifier (or VPI/VCI). A related  
10 aspect of invention encompasses software for implementing the automated provisioning of services through the access data network.

Fig. 7 is an enhanced block diagram of the access data network, showing certain additional elements involved in the automated selection and provisioning of customer services. Elements of the network referenced by the same or similar  
15 numbers as shown in the earlier drawings are substantially the same as like numbered elements discussed above, for example relative to Figs. 1 and 2.

The methodology enables a customer that is capable of receiving vertical services to use a web-based software application to order service. The automated interaction of the user through the common web-based software application  
20 causes communications to software entities extant in one or more of the services domains, which coordinate the processes that will automatically provision the corresponding data circuit, according to both customer-selected parameters and detected network parameters. This service selection and provisioning technique significantly reduces, or completely removes, the errors present in the previous

- 49 -

For example, the illustrated embodiment shows a DHCP server 81 within the network services domain 33 (carrier's private portion associated with the vertical services domain through the node 18). The carrier will also operate a web server 85 from the vertical services domain 13, although this server could operate from the public Internet. The carrier also operates one or more provisioning servers 87, for example, via the network services domain 33. The physical network of the vertical services domain 13 and the network services domain 33 may overlap or be substantially the same. The network services domain, for example, may comprise a private virtual network set up for the carrier's own network-operations through the same network as forms the vertical services domain.

The user will run a standard Internet browser application 83, for example on a PC 25 or other appropriate data communication device. The web server 85 and the provisioning server 87 run one or more applications 91, 93 and 95 related to the service selection and provisioning operations, as discussed in more detail below.

Certain communications, discussed layer, may be exchanged between the servers, 81, 85 and 87, for example via the node 18 and physical links of the domains 33 and 13. The provisioning server 87 also communicates with the actual service transport elements of the access data network, such as the DSLAM 17, the switch 19, and the gateway router 29, to conduct a metering operation and/or to activate and control provisioning of services through such elements. The links between the provisioning server 87 and these elements, shown by dot-dash lines in the drawing, may go through secure logical channels through the network services domain 33 and the node 18 or may use other private data communication links.

The preferred embodiment allows the user to select services via a standard web browser interface, that is to say using a browser application 83. The browser application 83 communicates with a Service Application 91. The carrier also utilizes two software modules, a Metering Application 93 and a Provisioning Application 95 in the presently preferred embodiment. The Service Application 91 and the Metering Application 93 may run as one or more programs on the web server 85, the provisioning server 87 or some other server securely maintained by

- 50 -

the carrier in one of the vertically inserted domains. For example, the Service Application 91 may run as one or more of the programs on the web server 85; whereas the Metering Application 93 and the Provisioning Application 95 preferably run on one or more provisioning servers 87 in the network services domain 33. Alternatively, the Metering Application 93 could run on one of the actual service transport elements of the access data network, such as the DSLAM 17 or the switch 19.

The hardware of a server system, such as the server 85 or the server 87, corresponds to that of a typical general-purpose computer, comprising a central processing unit (CPU) formed for example of one or more microprocessors, a number of memory devices and an interface to the data communication network, in this case to the IP network forming the respective domain 13 or 33. Such a server computer may also provide a local or remote graphical user interface (GUI) for operation and control of the server by carrier personnel. Such an interface, for example, may comprise a common type of display, a keyboard and one or more of the common types of cursor controls. Various media, readable by such a computer system, may store or carry the executable code and any associated data for the respective Application, for example the web pages provided by the server 85 and/or the associated code of the Service Application 91, for the Metering Application 93 or for the Provisioning Application 95. Examples of such media include semiconductor and disk type memories, digital tapes, and the like. Computer readable media used by such systems also include various types of signals sent and received by computer systems for loading software code, associated data and/or web pages into the memory and/or the CPU of the respective server system hardware.

Figure 8 shows the actors, activities, and attendant messages involved in this exemplary method. Each activity resides in the domain of a particular actor: the customer domain 800, the carrier/PSTN domain 810, or the service domain 820.

With reference to Figs. 7 and 8, the method begins when the customer contacts the carrier and requests ADSL service (step 801). The carrier connects the ADSL circuit (step 811), which entails moving the customer's POTS circuit to

- 51 -

the ADSL frame at the remote terminal or central office, and cross-connecting back to the switch frame.

In step 802, after receiving and setting up the ADSL equipment and concomitant software, the customer starts an IP session over the ADSL line. In 5 doing so, the customer's PC 25 (acting as a DHCP client) automatically requests and negotiates an IP address (step 8211) and other operational parameters with a DHCP server 81 in the network services domain 33. This IP address obtained in step 8211 is temporary, as it is used solely for the purpose of establishing a session. Later, the customer will receive a permanent or session IP address from 10 its ISP/Global Service Provider (GSP) 11'.

In step 803, the customer uses the Web browser application 83 to access one or more of the network service provider's service pages via the PC 25, for example, from the web server 85 on the vertical services domain 13. In order to avail the customer of the services and GSPs 11' that will work acceptably well on 15 the customer's particular ADSL line 21, the Service Application 91, determines the available services by acquiring a bandwidth metric (e.g. packet rate) from the Metering Application 93 (step 822). As shown in steps 823 and 804, the Metering Application 93 employs an interactive method for determining the maximum data rate. This rate determining method may be based on a damped 20 oscillation algorithm, and may employ the mechanisms inherent to TCP. In the preferred embodiment, the Metering Application 93 should be in the core of the service domain 13 or 33 so as to measure the TCP/IP packet rate between the subscriber equipment and the respective services domain. However, other metering application techniques may be used from the same or other locations.

25 In step 824, the Service Application 91 sends a Web page to the customer, populated with the available services and identifications of the GSPs. In step 805, the customer selects services and one (1) primary GSP. Step 824 involves a determination of the customer's data rate, which will be the lower of the metered rate and a rate based on the customer-selected services.

30 In step 825, the Service Application 91 receives the customer's selections, and requests the data circuit by sending customer's IP address and data rate. In

- 52 -

step 826, the Provisioning Application 95 obtains the customer's MAC address from the DHCP server 81.

In step 827, the Provisioning Application establishes communications with the vertical services insertion (VSI) ATM switch 19, and may use a script or other 5 remote mechanism to either request the appropriate data circuit, or set up the data circuit (8121). The data circuit will have an association with the customer's MAC address, and will have a provisioned peak rate no less than that initially determined by the Service Application 91.

In step 828, the Provisioning Application 95 receives the logical circuit 10 information (or VCI/VPI) for the customer, and subsequently communicates with the gateway router 29 to set up the appropriate logical circuit (or VPI/VCI) for the customer (step 8131). In step 813, the gateway router 29 may set up the circuit in response to a request from the Provisioning Application, or may be controlled by the Provisioning Application. The Provisioning Application 95 may 15 similarly communicate with the switch 19 or the DSLAM 17, as appropriate, to provision the customer circuit. Thus, the customer has a PVC that traverses the exemplary ADSL data network, and this PVC was set up using the automated method presented here.

Certain aspects of the invention relate to the software elements, such as 20 the combination of the Service Application 91, the Metering Application 93 and the Provisioning Application 95, shown in Fig. 7. At different times all or portions of the executable code for any or all of these elements may reside in physical media or be carried by electromagnetic media. Physical media include the memory of the server(s) 85, 87, such as various semiconductor memories, 25 tape drives, disc drives and the like of general-purpose computer systems. All or portions of the software may at times be communicated through various networks, for example to load the software from another computer into the respective server or into another network element. Thus, another type of media that may bear the software includes optical, electrical and electromagnetic waves, 30 such as used across physical interfaces between local devices, through wired and optical landline networks and over various air-links.

- 53 -

While the foregoing has described what are considered to be the best mode and/or other preferred embodiments of the invention, it is understood that various modifications may be made therein and that the invention may be implemented in various forms and embodiments, and that it may be applied in numerous applications, only some of which have been described herein. It is intended by the following claims to claim any and all modifications and variations that fall within the true scope of the inventive concepts.

5

**WHAT IS CLAIMED IS:**

1. In an access data network, comprising:
  - a communication access node coupled to a wide area internetwork;
  - 5 a plurality of digital subscriber line transceivers coupled to network ends of subscriber lines, for data communication with transceivers coupled to respective customer premises ends of respective subscriber lines;
  - 10 an access switch coupled for data communications with the digital subscriber line transceivers, for receiving data from customer premises equipment via respective ones of the digital subscriber line transceivers and for supplying data intended for transmission to predetermined customer premises equipment to the respective ones of the digital subscriber line transceivers;
  - 15 a high-speed data link between the access switch and the communication access node;
  - 20 a vertical services network coupled locally to the access switch, wherein the access switch segregates upstream traffic from the respective customer premises equipment between the vertical services network and the high-speed data link to the communication access node, and the access switch aggregates downstream traffic from the vertical services network and the high-speed data link from the communication access node for transmission to the respective customer premises equipment,
- 25 a method of establishing services for one of the customers in said access data network, comprising the steps of:
  - conducting an interactive session between customer premises equipment of the one customer and a server coupled to the vertical services network, to obtain customer service selection data from the one customer;
  - 30 automatically determining a capability of the access data network to provide a service to the customer premises equipment of the one customer; and
  - automatically provisioning a logical circuit for the one customer extending from the customer premises equipment of the one customer to the access node

via the access switch and the high-speed data link, in accord with an appropriate one or more of the determined capability and the customer service selection data.

2. A method as in claim 1, wherein the logical circuit comprises a layer-2 protocol logical communication circuit provisioned through the access  
5 switch and the high-speed data link.

3. A method as in claim 2, wherein the layer-2 protocol logical communication circuit comprises an Asynchronous Transfer Mode (ATM) virtual circuit.

4. A method as in claim 1, further comprising automatically  
10 provisioning at least one additional service provided from the vertical services network in response to at least a portion of the customer service selection data.

5. A method as in claim 1, wherein the provisioning of the logical circuit comprises setting a peak rate for the logical service that is no less than: (1) a metered maximum rate for service to the customer premises equipment of the  
15 one customer found in the determining step, and (2) a rate associated with a customer selected service identified in the customer service selection data.

6. A communication system with automated service provisioning capabilities, comprising:

an access data network, for providing access services to at least two  
20 different network domains; and

a web server, coupled to a first one of the domains of the access data network, for interactive communication with a customer subscribing to services;  
and

means responsive to customer service selection data obtained via the  
25 interactive communication with the web server, for automatically determining capabilities of the access data network to provide service for the customer and for

- 56 -

automatically provisioning at least a virtual circuit through the access data network to a second one of the domains in accord with an appropriate one or more of the determined capabilities and the customer service selection data.

7. A communication system as in claim 6, said means provisions the  
5 virtual circuit to have a peak allowable rate at or above the lower of: (1) a rate corresponding to a customer-selected service identified in the customer service selection data, and a metered maximum rate for service to equipment of the customer found by the automatic determination of the capabilities of the access data network.

10 8. A software product enabling automated service provisioning of an access data communication network, the software product comprising:

at least one machine readable medium;  
programming code, carried by the at least one machine readable medium,  
for execution by at least one computer, wherein the programming code  
15 comprises:

a services application for implementing interactive communication with a customer subscribing to services of the access data communication network to obtain service selection information from the customer;

20 a metering application for automatically determining capabilities of the access data communication network to provide service for the customer; and

a provisioning application for provisioning a logical circuit through the access data communication network in accord with an appropriate one or more of the service selection information and the determined capabilities.

9. A software product as in claim 6, wherein:

25 the service selection information includes a customer-selected service having an associated rate of service through the access data communication network;

- 57 -

the determined capabilities comprise a metered maximum rate of data communication that may be supported by the access data communication network to equipment of the customer; and

5       the provisioning application provisions the logical circuit to allow a peak rate no less than the lower of the rate of service associated with the a customer-selected service and the metered maximum rate.

10.      A method for automated service provisioning in an access data communication network, comprising the steps of:

10       interacting through the access data communication network to obtain service selection information from a customer;

automatically determining capabilities of the access data communication network to provide service for the customer; and

15       automatically provisioning a logical circuit through the access data communication network in accord with an appropriate one or more of the determined capabilities and the customer service selection data.

1/10

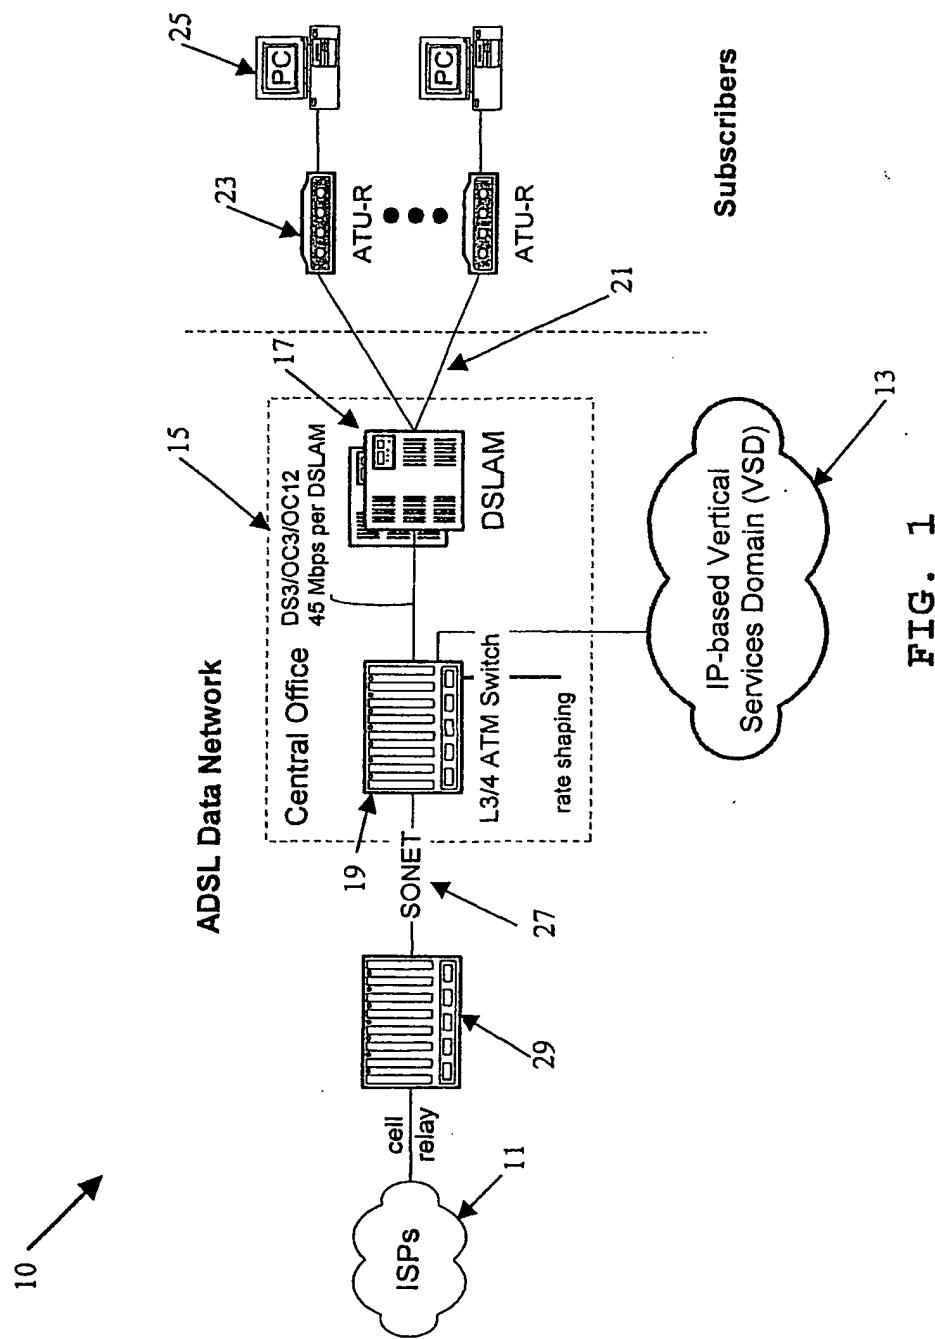


FIG. 1

2/10

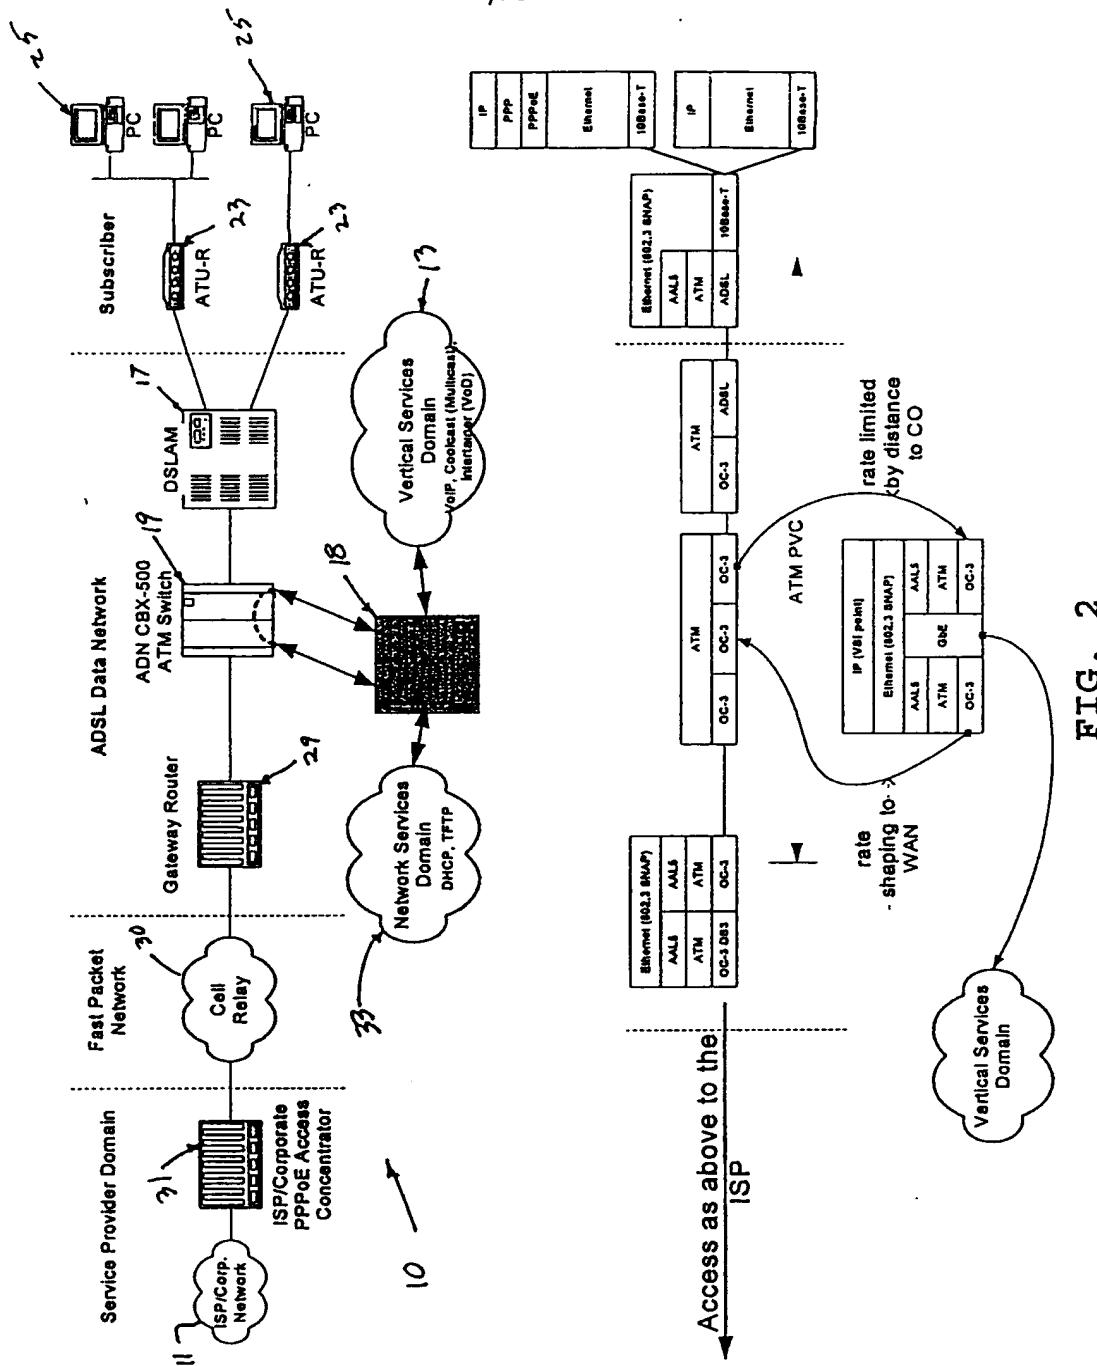


FIG. 2

3/10

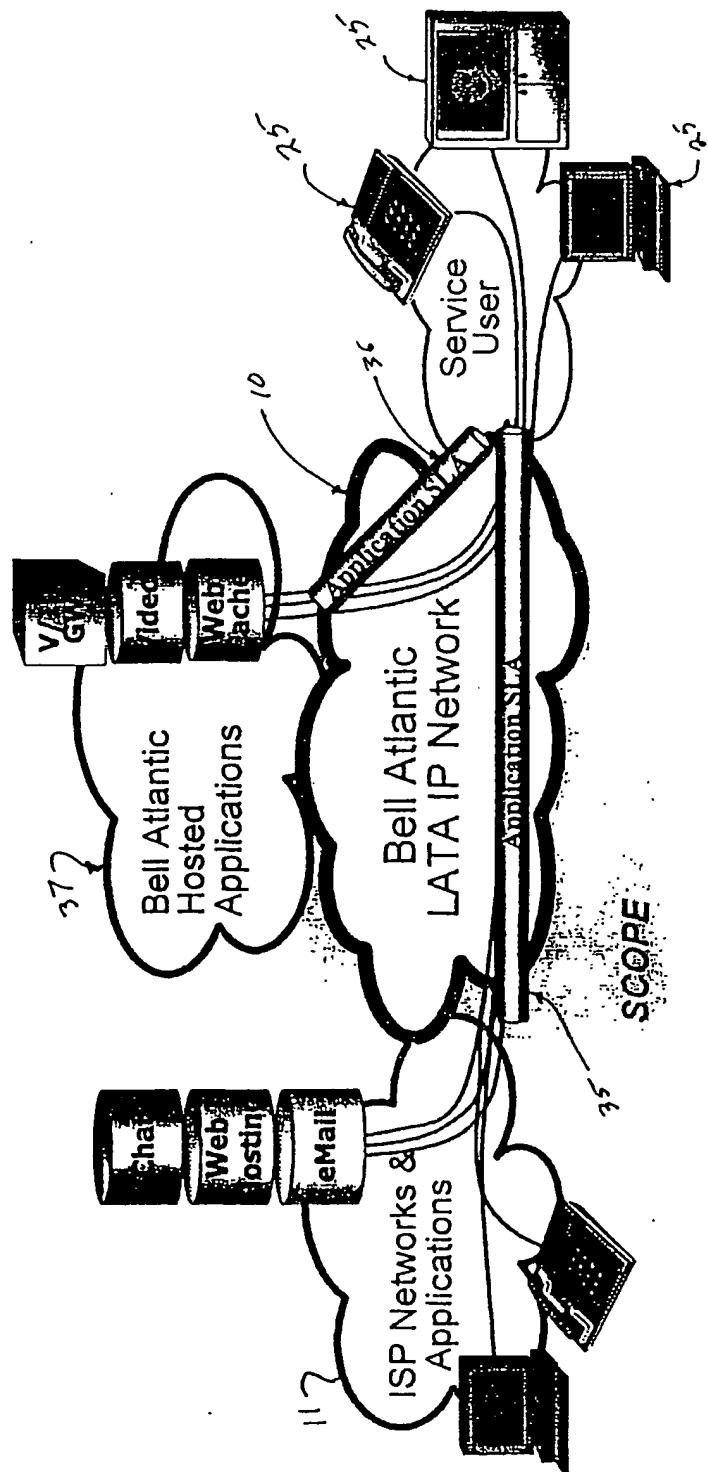


FIG. 3

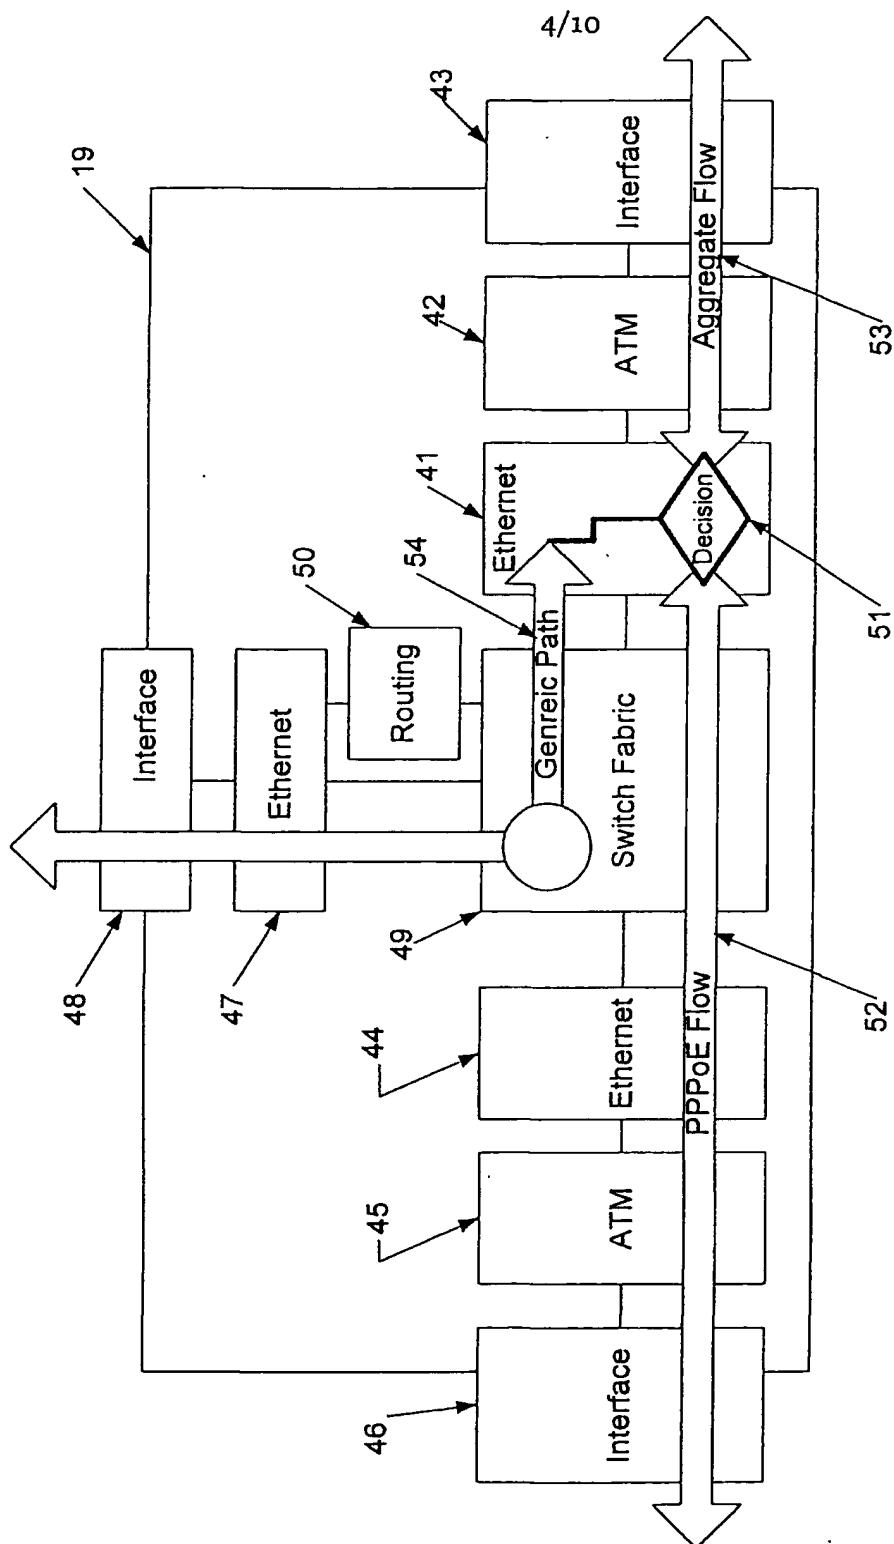


FIG. 4

5/10

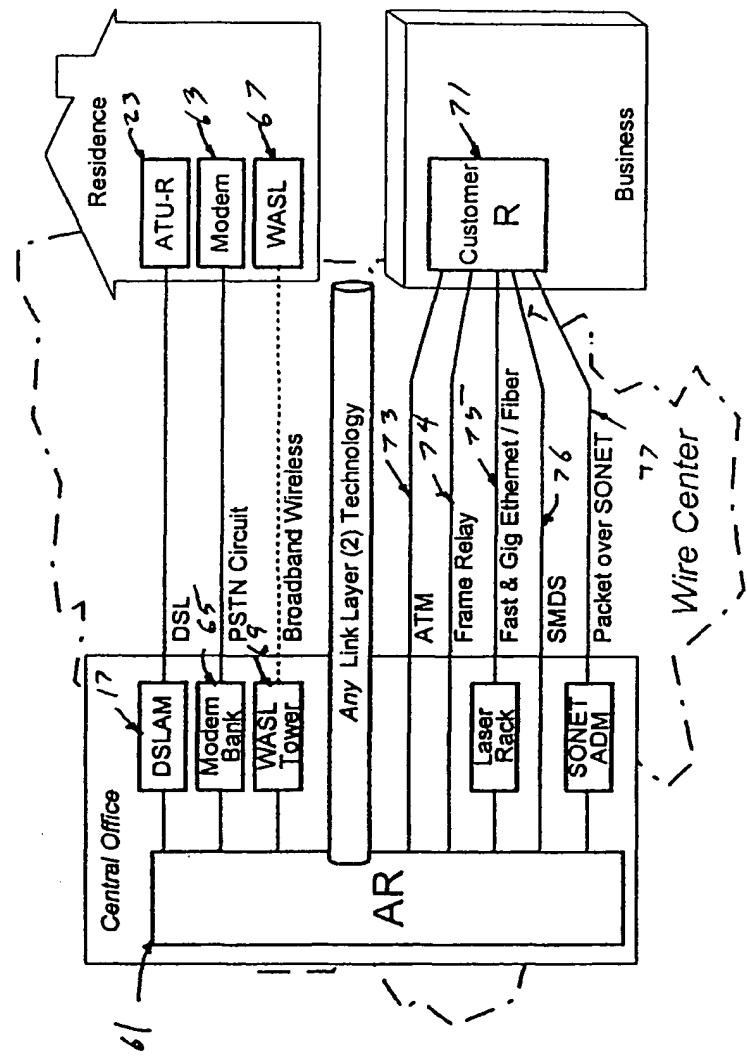


FIG. 5

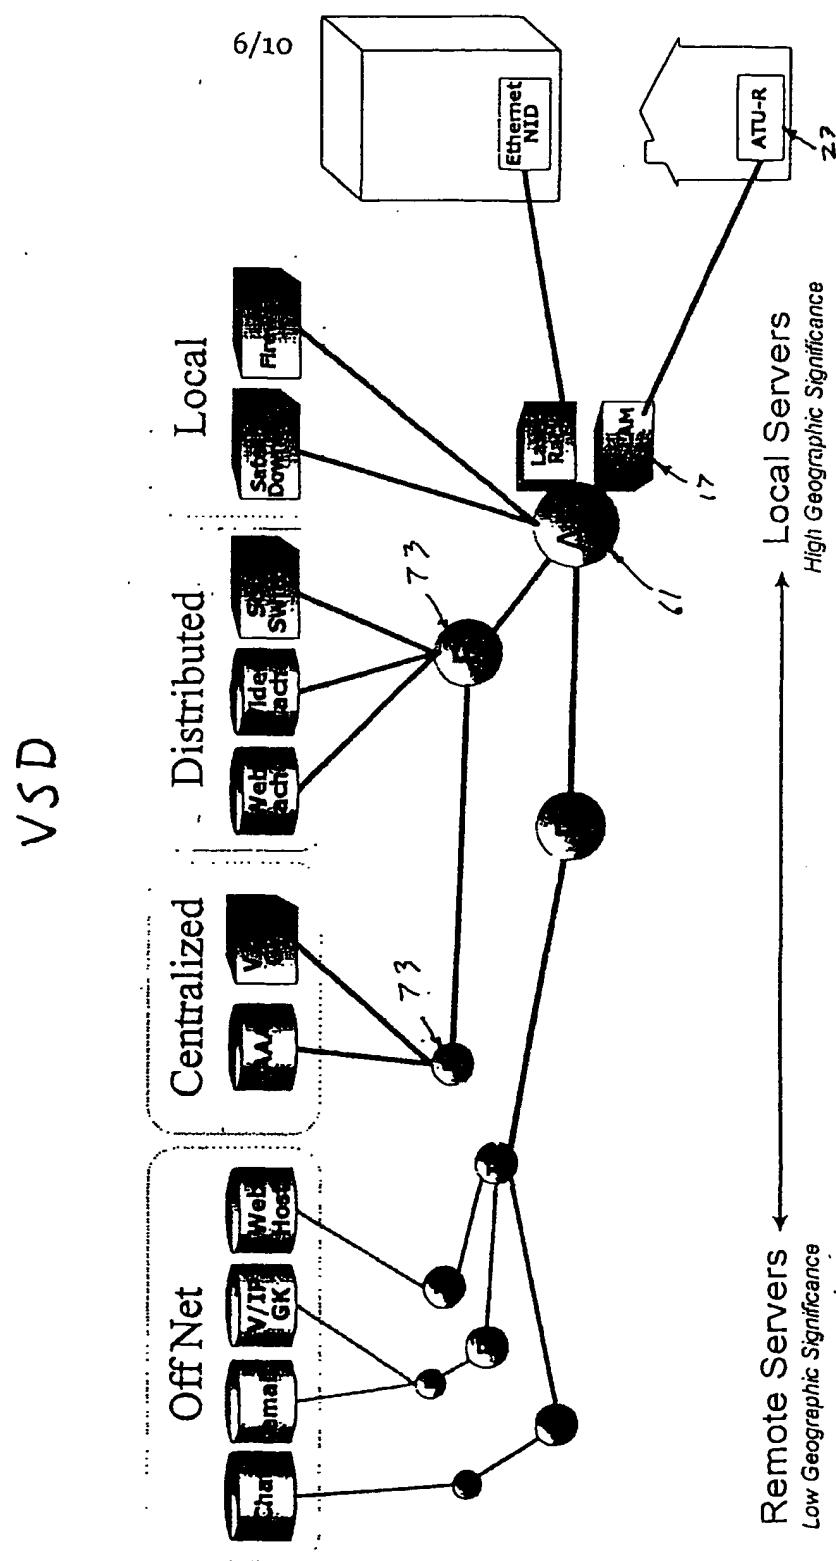


FIG. 6

7/10

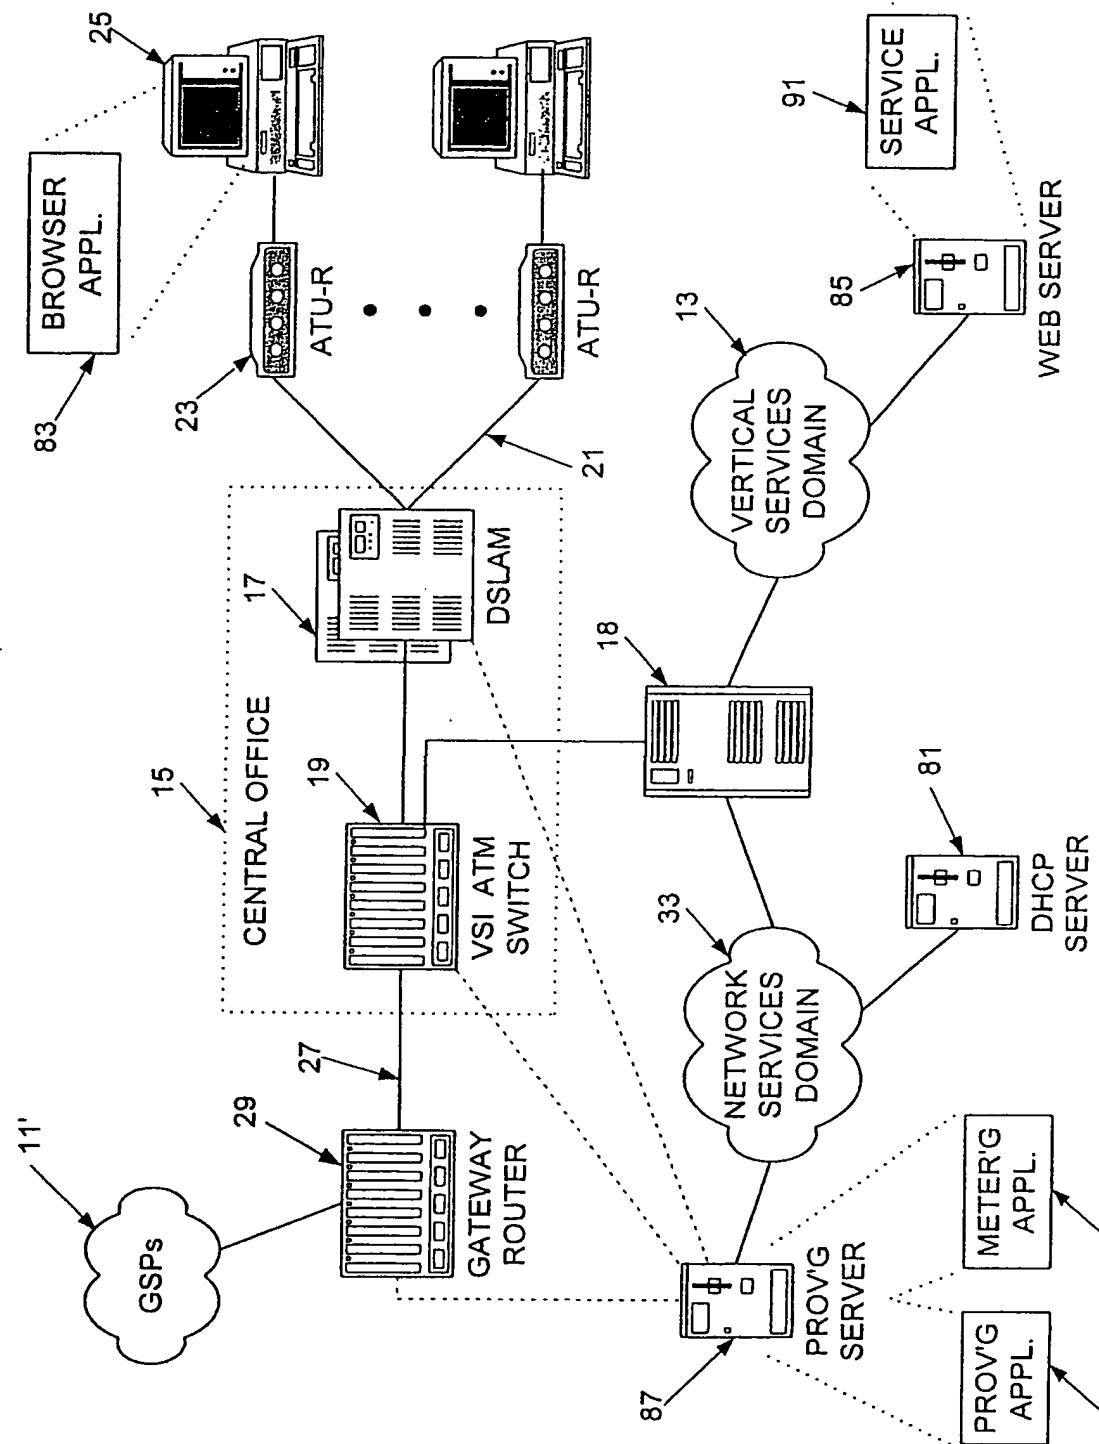


FIG. 7

8/10

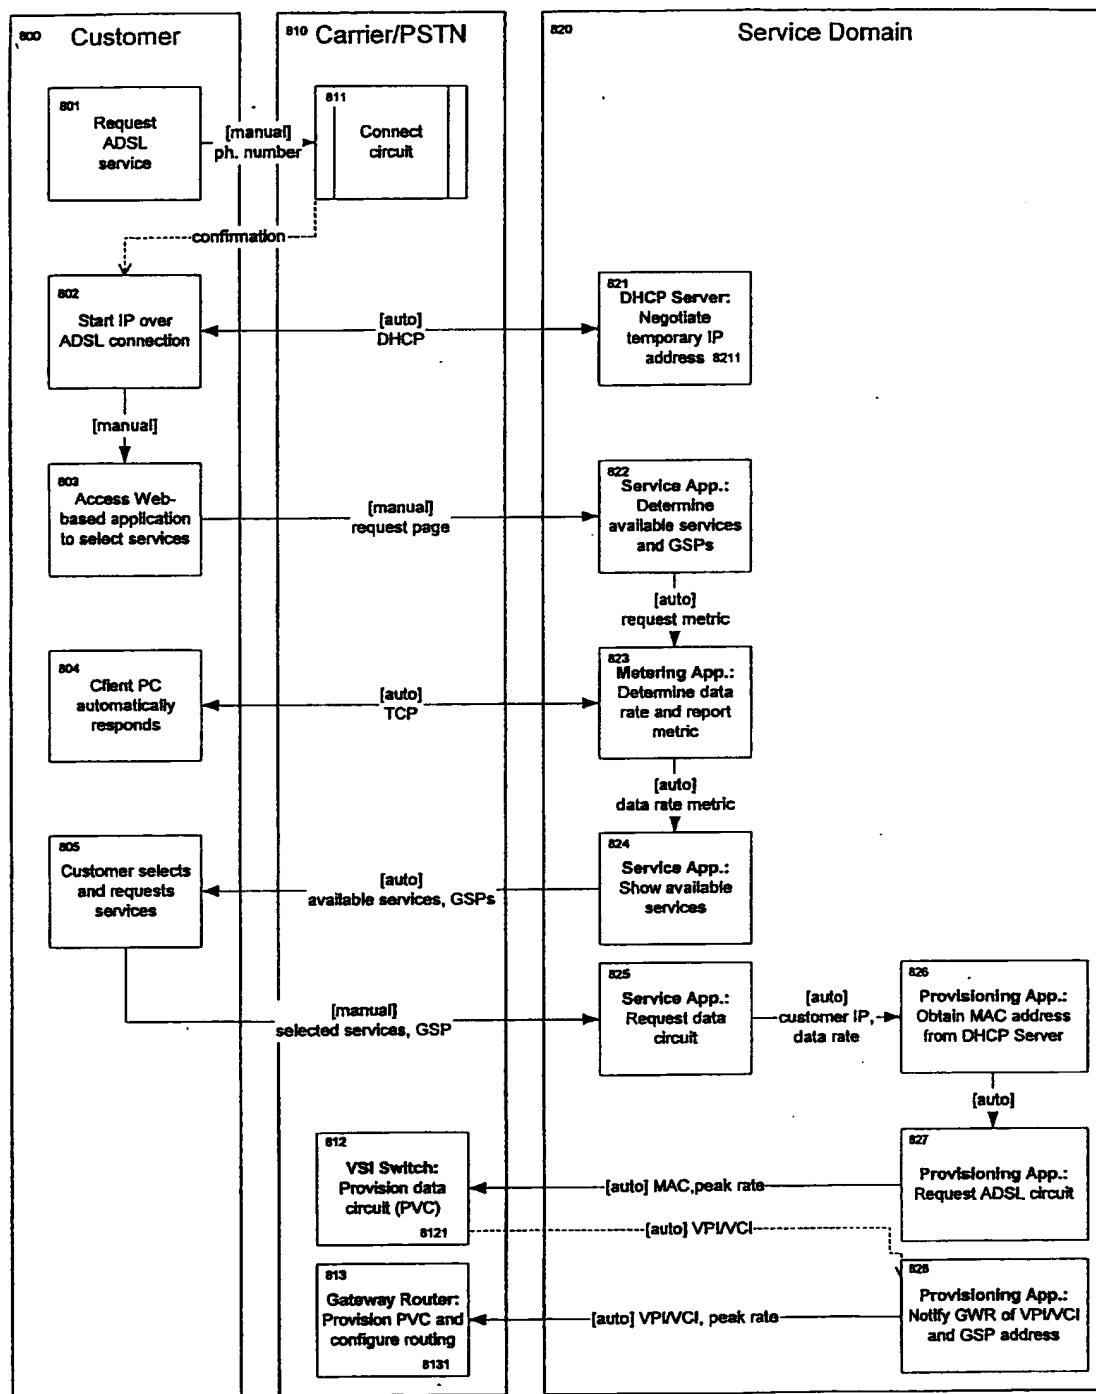
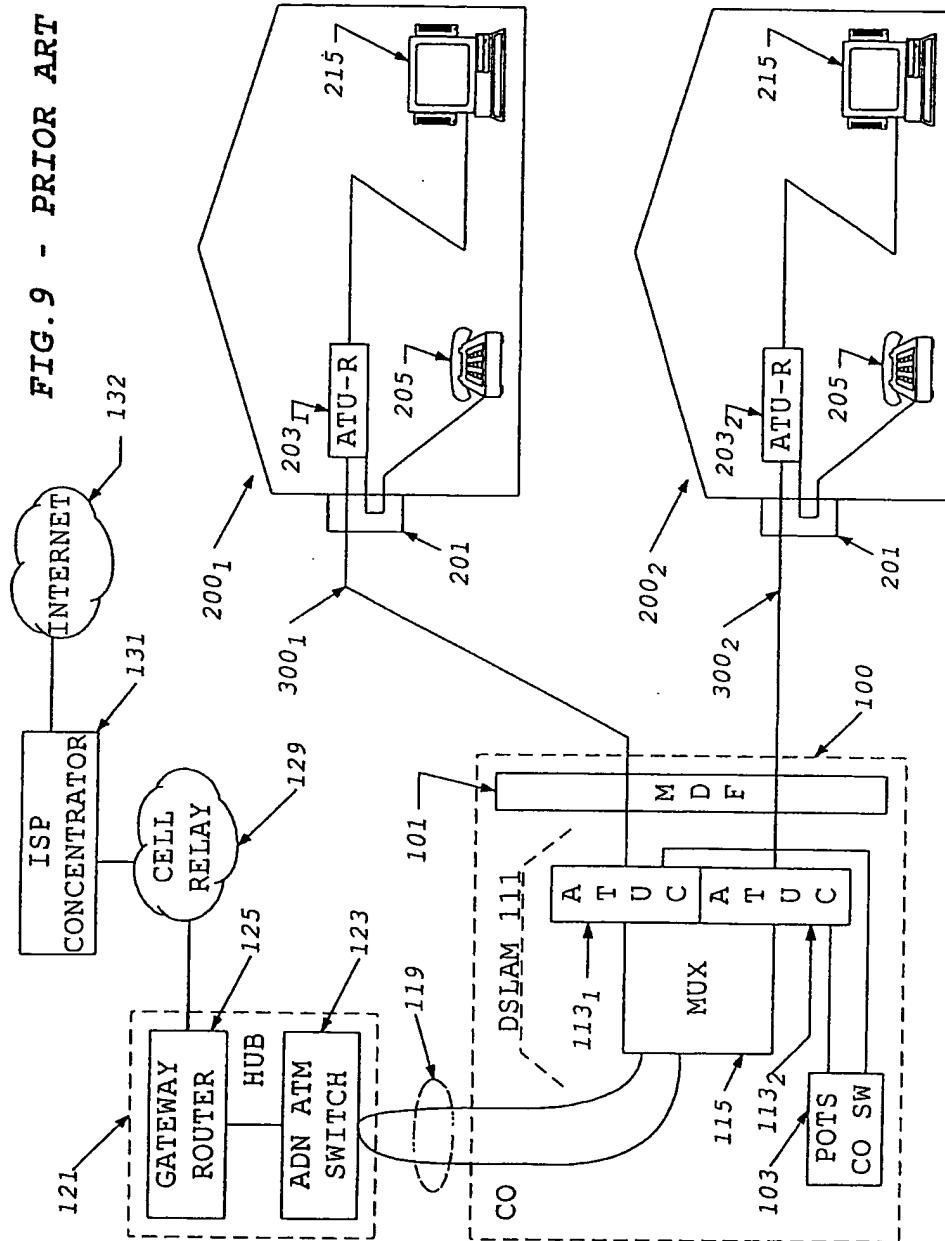


FIG. 8

9/10

FIG. 9 - PRIOR ART



10/10

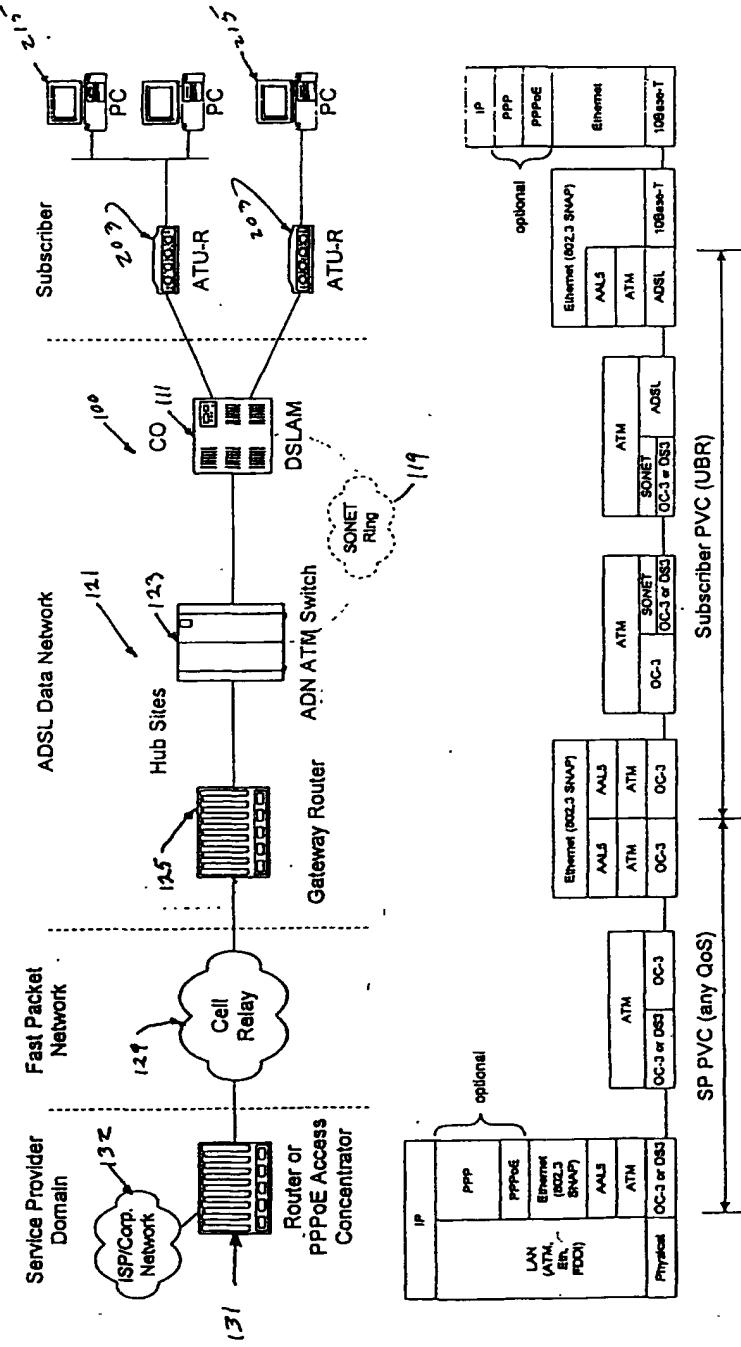


FIG. 10 - PRIOR ART

## INTERNATIONAL SEARCH REPORT

Int'l application No.  
PCT/US01/24855**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(7) :H04L 12/56; H04J 3/22, 3/24

US CL :370/394, 395, 397, 399, 401, 409, 412, 465, 466, 467, 470, 474

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/394, 395, 397, 399, 401, 409, 412, 465, 466, 467, 470, 474

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EAST Database

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,790,548 A (SISTANIZADEH et al) 04 August 1998, see the entire document.	1, 4-7, 10
A	US 5,905,781 A (MCHALE et al) 18 May 1999, see the entire document.	2, 3, 8, 9
X,P	US 6,118,785 A (ARAUJO et al) 12 September 2000, see the entire document.	2, 3, 8, 9

<input type="checkbox"/>	Further documents are listed in the continuation of Box C.	<input type="checkbox"/>	See patent family annex.
"A"	Special categories of cited documents document defining the general state of the art which is not considered to be of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"&"	document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search  17 SEPTEMBER 2001	Date of mailing of the international search report  19 NOV 2001
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-9230	Authorized officer NGUYEN TOAN Telephone No. (703) 305-0140

**THIS PAGE BLANK (USPTO)**